Under consideration for publication in J. Functional Programming

Pushdown flow analysis with abstract garbage collection

J. IAN JOHNSON
Northeastern University

ILYA SERGEY
IMDEA Software Institute

CHRISTOPHER EARL
University of Utah

MATTHEW MIGHT
University of Utah

DAVID VAN HORN
University of Maryland

arXiv:1406.5106v1 [cs.PL] 19 Jun 2014

Abstract

In the static analysis of functional programs, pushdown flow analysis and abstract garbage collection push the boundaries of what we can learn about programs statically. This work illuminates and poses solutions to theoretical and practical challenges that stand in the way of combining the power of these techniques. Pushdown flow analysis grants unbounded yet computable polyvariance to the analysis of return-flow in higher-order programs. Abstract garbage collection grants unbounded polyvariance to abstract addresses which become unreachable between invocations of the abstract contexts in which they were created. Pushdown analysis solves the problem of precisely analyzing recursion in higher-order languages; abstract garbage collection is essential in solving the "stickiness" problem. Alone, our benchmarks demonstrate that each method can reduce analysis times and boost precision by orders of magnitude.

We combine these methods. The challenge in marrying these techniques is not subtle: computing the reachable control states of a pushdown system relies on limiting access during transition to the top of the stack; abstract garbage collection, on the other hand, needs full access to the entire stack to compute a root set, just as concrete collection does. *Conditional* pushdown systems were developed for just such a conundrum, but existing methods are ill-suited for the dynamic nature of garbage collection.

We show fully precise and approximate solutions to the feasible paths problem for pushdown garbage-collecting control-flow analysis. Experiments reveal synergistic interplay between garbage collection and pushdown techniques, and the fusion demonstrates "better-than-both-worlds" precision.

```
(define (id x) x)

(define (f n)
  (cond [(<= n 1)  1]
        [else      (* n (f (- n 1)))]))

(define (g n)
  (cond [(<= n 1)  1]
        [else      (+ (* n n) (g (- n 1)))]))

(print (+ ((id f) 3) ((id g) 4)))
```

Fig. 1: A small example to illuminate the strengths and weaknesses of both pushdown analysis and abstract garbage collection.

1 Introduction

The development of a context-free[1] approach to control-flow analysis (CFA2) by Vardoulakis and Shivers (2010) provoked a shift in the static analysis of higher-order programs. Prior to CFA2, a precise analysis of recursive behavior had been a challenge—even though flow analyses have an important role to play in optimization for functional languages, such as flow-driven inlining (Might and Shivers 2006a), interprocedural constant propagation (Shivers 1991) and type-check elimination (Wright and Jagannathan 1998).

While it had been possible to statically analyze recursion *soundly*, CFA2 made it possible to analyze recursion *precisely* by matching calls and returns without approximating the stack as k-CFA does. The approximation is only in the binding structure, and not the control structure of the program. In its pursuit of recursion, clever engineering steered CFA2 to a *theoretically* intractable complexity, though in practice it performs well. Its payoff is significant reductions in analysis time *as a result of* corresponding increases in precision.

For a visual measure of the impact, Figure 2 renders the abstract transition graph (a model of all possible traces through the program) for the toy program in Figure 1. For this example, pushdown analysis eliminates spurious return-flow from the use of recursion. But, recursion is just one problem of many for flow analysis. For instance, pushdown analysis still gets tripped up by the spurious cross-flow problem; at calls to `(id f)` and `(id g)` in the previous example, it thinks `(id g)` could be `f` *or* `g`. CFA2 is not confused in this due to its precise stack frames, but can be confused by unreachable heap-allocated bindings.

Powerful techniques such as abstract garbage collection (Might and Shivers 2006b) were developed to address the cross-flow problem (here in a way complementary to CFA2's stack frames). The cross-flow problem arises because monotonicity prevents revoking a judgment like "procedure `f` flows to `x`," or "procedure `g` flows to `x`," once it's been made.

[1] As in context-free language, not a context-insensitive analysis.

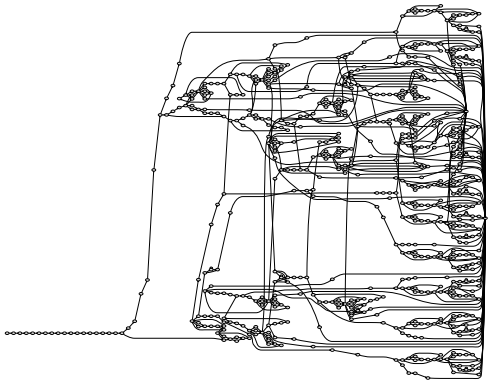

(1) without pushdown analysis or abstract GC: 653 states

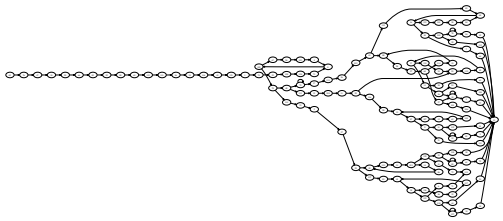

(2) with pushdown only: 139 states

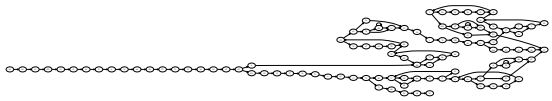

(3) with GC only: 105 states

(4) with pushdown analysis and abstract GC: 77 states

Fig. 2: We generated an abstract transition graph for the same program from Figure 1 four times: (1) without pushdown analysis or abstract garbage collection; (2) with only abstract garbage collection; (3) with only pushdown analysis; (4) with both pushdown analysis and abstract garbage collection. With only pushdown or abstract GC, the abstract transition graph shrinks by an order of magnitude, but in different ways. The pushdown-only analysis is confused by variables that are bound to several different higher-order functions, but for short durations. The abstract-GC-only is confused by non-tail-recursive loop structure. With both techniques enabled, the graph shrinks by nearly half yet again and fully recovers the control structure of the original program.

In fact, abstract garbage collection, by itself, also delivers significant improvements to analytic speed and precision in many benchmarks. (See Figure 2 again for a visualization of that impact.)

It is natural to ask: can abstract garbage collection and pushdown analysis work together? Can their strengths be multiplied? At first, the answer appears to be a disheartening "*No*."

1.1 The problem: The whole stack versus *just the top*

Abstract garbage collection seems to require more than pushdown analysis can decidably provide: access to the full stack. Abstract garbage collection, like its name implies, discards unreachable values from an abstract store during the analysis. Like concrete garbage collection, abstract garbage collection also begins its sweep with a root set, and like concrete garbage collection, it must traverse the abstract stack to compute that root set. But, pushdown systems are restricted to viewing the top of the stack (or a bounded depth)—a condition violated by this traversal.

Fortunately, abstract garbage collection does not need to arbitrarily modify the stack. It only needs to know the root set of addresses in the stack. This kind of system has been studied before in the context of compilers that build a symbol table (a so-called "one-way stack automaton" (Ginsburg et al. 1967)),in the context of first-order model-checking (pushdown systems with checkpoints (Esparza et al. 2003)),and also in the context of points-to analysis for Java (conditional weighted pushdown systems (CWPDS) (Li and Ogawa 2010)). We borrow the definition of (unweighted) conditional pushdown system (CPDS) in this work, though our analysis does not take CPDSs as inputs.

Higher-order flow analyses typically do not take a control-flow graph, or similar pre-abstracted object, as input and produce an annotated graph as output. Instead, they take a program as input and "run it on all possible inputs" (abstractly) to build an approximation of the language's reduction relation (semantics), specialized to the given program. This semantics may be non-standard in such a way that extra-semantic information might be accumulated for later analyses' consumption. The important distinction between higher-order and first-order analyses is that the *model* to analyze is built *during* the analysis, which involves interpreting the program (abstractly).

When a language's semantics treats the control stack as an actual stack, *i.e.*, it does not have features such as first-class continuations, an interpreter can be split into two parts: a function that takes the current state and returns all next states along with a pushed activation frame or a marker that the stack is unchanged; and a function that takes the current state, a possible "top frame" of the stack, and returns the next states after popping this frame. This separation is crucial for an effective algorithm, since pushed frames are understood from program text, and popped frames need only be enumerated from a (usually small) set that we compute along the way.

Control-state reachability for the straightforward formulation of stack introspection ends up being uncomputable. Conditional pushdown systems introduce a relatively weak regularity constraint on transitions' introspection: a CPDS may match the current stack against a choice of finitely many regular languages of stacks in order to transition from one state to the next along with the stack action. The general solutions to feasible paths in *conditional*

pushdown systems enumerate all languages of stacks that a transition may be conditioned on. This strategy is a non-starter for garbage collection, since we delineate stacks by the addresses they keep live; this is exponential in the number of addresses. The abstraction step that finitizes the address space is what makes the problem fall within the realm of CPDSs, even if the target is so big it barely fits. But abstract garbage collection is special — we can compute which languages of stacks we need to check against, given the current state of the analysis. It is therefore possible to fuse the full benefits of abstract garbage collection with pushdown analysis. The dramatic reduction in abstract transition graph size from the top to the bottom in Figure 2 (and echoed by later benchmarks) conveys the impact of this fusion.

Secondary motivations There are four secondary motivations for this work:

1. bringing context-sensitivity to pushdown analysis;
2. exposing the context-freedom of the analysis;
3. enabling pushdown analysis without continuation-passing style; and
4. defining an alternative algorithm for computing pushdown analysis, introspectively or otherwise.

In CFA2, monovariant (0CFA-like) context-sensitivity is etched directly into the abstract "local" semantics, which is in turn phrased in terms of an explicit (imperative) summarization algorithm for a partitioned continuation-passing style. Our development exposes the classical parameters (exposed as allocation functions in a semantics) that allow one to tune the context-sensitivity and polyvariance (accomplishing (1)), thanks to the semantics of the analysys formulated in the form of an "abstracted abstract machine" (Van Horn and Might 2012).

In addition, the context-freedom of CFA2 is buried implicitly inside an imperative summarization algorithm. No pushdown system or context-free grammar is explicitly identified. Thus, a motivating factor for our work was to make the pushdown system in CFA2 explicit, and to make the control-state reachability algorithm purely functional (accomplishing (2)).

A third motivation was to show that a transformation to continuation-passing style is unnecessary for pushdown analysis. In fact, pushdown analysis is arguably more natural over direct-style programs. By abstracting all machine components except for the program stack, it converts naturally and readily into a pushdown system (accomplishing (3)). In his dissertation, Vardoulakis showed a direct-style version of CFA2 that exploits the meta-language's runtime stack to get precise call-return matching. The approach is promising, but its correctness remains unproven, and it does not apply to generic pushdown systems.

Finally, to bring much-needed clarity to algorithmic formulation of pushdown analysis, we have included an appendix containing a reference implementation in Haskell (accomplishing (4)). We have kept the code as close in form to the mathematics as possible, so that where concessions are made to the implementation, they are obvious.

1.2 Overview

We first review preliminaries to set a consistent feel for terminology and notation, particularly with respect to pushdown systems. The derivation of the analysis begins with a con-

crete CESK-machine-style semantics for A-Normal Form λ-calculus. The next step is an infinite-state abstract interpretation, constructed by bounding the C(ontrol), E(nvironment) and S(tore) portions of the machine. Uncharacteristically, we leave the stack component—the K(ontinuation)—unbounded.

A shift in perspective reveals that this abstract interpretation is a pushdown system. We encode it as a pushdown automaton explicitly, and pose control state reachability as a decidable language intersection problem. We then extract a rooted pushdown system from the pushdown automaton. For completeness, we fully develop pushdown analysis for higher-order programs, including an efficient algorithm for computing reachable control states. We go further by characterizing complexity and demonstrating the approximations necessary to get to a polynomial-time algorithm.

We then introduce abstract garbage collection and quickly find that it violates the pushdown model with its traversals of the stack. To prove the decidability of control-state reachability, we formulate introspective pushdown systems, and recast abstract garbage collection within this framework. We then review that control-state reachability is decidable for introspective pushdown systems as well when subjected to a straightforward regularity constraint.

We conclude with an implementation and empirical evaluation that shows strong synergies between pushdown analysis and abstract garbage collection, including significant reductions in the size of the abstract state transition graph.

1.3 Contributions

We make the following contributions:

1. Our primary contribution is an *online* decision procedure for reachability in introspective pushdown systems, with a more efficient specialization to abstract garbage collection.
2. We show that classical notions of context-sensitivity, such as *k*-CFA and poly/CFA, have direct generalizations in a pushdown setting. CFA2 was presented as a monovariant analysis,[2] whereas we show polyvariance is a natural extension.
3. We make the context-free aspect of CFA2 explicit: we clearly define and identify the pushdown system. We do so by starting with a classical CESK machine and systematically abstracting until a pushdown system emerges. We also remove the orthogonal frame-local-bindings aspect of CFA2, so as to focus solely on the pushdown nature of the analysis.
4. (*) We remove the requirement for a global CPS-conversion by synthesizing the analysis directly for direct-style (in the form of A-normal form lambda-calculus — a local transformation).
5. We empirically validate claims of improved precision on a suite of benchmarks. We find synergies between pushdown analysis and abstract garbage collection that makes the whole greater that the sum of its parts.

[2] Monovariance refers to an abstraction that groups all bindings to the same variable together: there is *one* abstract variant for all bindings to each variable.

6. We provide a mirror of the major formal development as working Haskell code in the appendix. This code illuminates dark corners of pushdown analysis and it provides a concise formal reference implementation.

(*) The CPS requirement distracts from the connection between continuations and stacks. We do not discuss `call/cc` in detail, since we believe there are no significant barriers to adapting the techniques of Vardoulakis and Shivers (2011) to the direct-style setting, given related work in Johnson and Van Horn (2013). Languages with exceptions fit within the pushdown model since a throw can be modeled as "pop until first catch."

2 Pushdown Preliminaries

The literature contains many equivalent definitions of pushdown machines, so we adapt our own definitions from Sipser (2005). *Readers familiar with pushdown theory may wish to skip ahead.*

2.1 Syntactic sugar

When a triple (x, ℓ, x') is an edge in a labeled graph:

$$x \overset{\ell}{\rightarrowtail} x' \equiv (x, \ell, x').$$

Similarly, when a pair (x, x') is a graph edge:

$$x \rightarrowtail x' \equiv (x, x').$$

We use both string and vector notation for sequences:

$$a_1 a_2 \ldots a_n \equiv \langle a_1, a_2, \ldots, a_n \rangle \equiv \vec{a}.$$

2.2 Stack actions, stack change and stack manipulation

Stacks are sequences over a stack alphabet Γ. To reason about stack manipulation concisely, we first turn stack alphabets into "stack-action" sets; each character represents a change to the stack: push, pop or no change.

For each character γ in a stack alphabet Γ, the **stack-action** set $\Gamma_\pm$ contains a push character γ_+; a pop character γ_-; and a no-stack-change indicator, ε:

$$\begin{array}{lll} g \in \Gamma_\pm ::= \varepsilon & & \text{[stack unchanged]} \\ \quad \mid \ \gamma_+ & \text{for each } \gamma \in \Gamma & \text{[pushed } \gamma\text{]} \\ \quad \mid \ \gamma_- & \text{for each } \gamma \in \Gamma & \text{[popped } \gamma\text{]}. \end{array}$$

In this paper, the symbol g represents some stack action.

When we develop introspective pushdown systems, we are going to need formalisms for easily manipulating stack-action strings and stacks. Given a string of stack actions, we can compact it into a minimal string describing net stack change. We do so through the operator $\lfloor \cdot \rfloor : \Gamma_\pm^* \to \Gamma_\pm^*$, which cancels out opposing adjacent push-pop stack actions:

$$\lfloor \vec{g}\, \gamma_+ \gamma_- \, \vec{g}\,' \rfloor = \lfloor \vec{g}\, \vec{g}\,' \rfloor \qquad\qquad \lfloor \vec{g}\, \varepsilon \, \vec{g}\,' \rfloor = \lfloor \vec{g}\, \vec{g}\,' \rfloor,$$

so that $\lfloor \vec{g} \rfloor = \vec{g}$, if there are no cancellations to be made in the string $\vec{g}$.

We can convert a net string back into a stack by stripping off the push symbols with the stackify operator, $\lceil \cdot \rceil : \Gamma_{\pm}^{*} \rightharpoonup \Gamma^{*}$:

$$\lceil \gamma_{+} \gamma'_{+} \dots \gamma_{+}^{(n)} \rceil = \langle \gamma^{(n)}, \dots, \gamma', \gamma \rangle,$$

and for convenience, $\lceil \vec{g} \rceil = \lceil \lfloor \vec{g} \rfloor \rceil$. Notice the stackify operator is defined for strings containing only push actions.

2.3 Pushdown systems

A **pushdown system** is a triple $M = (Q, \Gamma, \delta)$ where:

1. Q is a finite set of control states;
2. Γ is a stack alphabet; and
3. $\delta \subseteq Q \times \Gamma_{\pm} \times Q$ is a transition relation.

The set $Q \times \Gamma^{*}$ is called the **configuration-space** of this pushdown system. We use $\mathbb{PDS}$ to denote the class of all pushdown systems.

For the following definitions, let $M = (Q, \Gamma, \delta)$.

- The labeled **transition relation** $(\overset{}{\longmapsto}_{M}) \subseteq (Q \times \Gamma^{*}) \times \Gamma_{\pm} \times (Q \times \Gamma^{*})$ determines whether one configuration may transition to another while performing the given stack action:

$$(q, \vec{\gamma}) \xmapsto[M]{\varepsilon} (q', \vec{\gamma}) \text{ iff } q \overset{\varepsilon}{\rightarrowtail} q' \in \delta \qquad \text{[no change]}$$

$$(q, \gamma : \vec{\gamma}) \xmapsto[M]{\gamma_{-}} (q', \vec{\gamma}) \text{ iff } q \overset{\gamma_{-}}{\rightarrowtail} q' \in \delta \qquad \text{[pop]}$$

$$(q, \vec{\gamma}) \xmapsto[M]{\gamma_{+}} (q', \gamma : \vec{\gamma}) \text{ iff } q \overset{\gamma_{+}}{\rightarrowtail} q' \in \delta \qquad \text{[push].}$$

- If unlabelled, the transition relation $(\longmapsto)$ checks whether *any* stack action can enable the transition:

$$c \underset{M}{\longmapsto} c' \text{ iff } c \xmapsto[M]{g} c' \text{ for some stack action } g.$$

- For a string of stack actions $g_1 \dots g_n$:

$$c_0 \xmapsto[M]{g_1 \cdots g_n} c_n \text{ iff } c_0 \xmapsto[M]{g_1} c_1 \xmapsto[M]{g_2} \cdots \xmapsto[M]{g_{n-1}} c_{n-1} \xmapsto[M]{g_n} c_n,$$

  for some configurations $c_0, \dots, c_n$.
- For the transitive closure:

$$c \underset{M}{\longmapsto}^{*} c' \text{ iff } c \xmapsto[M]{\vec{g}} c' \text{ for some action string } \vec{g}.$$

Note Some texts define the transition relation δ so that $\delta \subseteq Q \times \Gamma \times Q \times \Gamma^{*}$. In these texts, $(q, \gamma, q', \vec{\gamma}) \in \delta$ means, "if in control state q while the character γ is on top, pop the stack, transition to control state q' and push $\vec{\gamma}$." Clearly, we can convert between these two representations by introducing extra control states to our representation when it needs to push multiple characters.

2.4 Rooted pushdown systems

A **rooted pushdown system** is a quadruple (Q,Γ,δ,q_0) in which (Q,Γ,δ) is a pushdown system and $q_0 \in Q$ is an initial (root) state. $\mathbb{RPDS}$ is the class of all rooted pushdown systems. For a rooted pushdown system $M = (Q,\Gamma,\delta,q_0)$, we define the **reachable-from-root transition relation**:

$$c \overset{g}{\underset{M}{\longmapsto\!\!\!\!\!\rightarrow}} c' \text{ iff } (q_0,\langle\rangle) \underset{M}{\longmapsto}{}^{*} c \text{ and } c \overset{g}{\underset{M}{\longmapsto}} c'.$$

In other words, the root-reachable transition relation also makes sure that the root control state can actually reach the transition.

We overload the root-reachable transition relation to operate on control states:

$$q \overset{g}{\underset{M}{\longmapsto\!\!\!\!\!\rightarrow}} q' \text{ iff } (q,\vec{\gamma}) \overset{g}{\underset{M}{\longmapsto\!\!\!\!\!\rightarrow}} (q',\vec{\gamma}\,') \text{ for some stacks } \vec{\gamma},\vec{\gamma}\,'.$$

For both root-reachable relations, if we elide the stack-action label, then, as in the un-rooted case, the transition holds if *there exists* some stack action that enables the transition:

$$q \underset{M}{\longmapsto\!\!\!\!\!\rightarrow} q' \text{ iff } q \overset{g}{\underset{M}{\longmapsto\!\!\!\!\!\rightarrow}} q' \text{ for some action } g.$$

2.5 Computing reachability in pushdown systems

A pushdown flow analysis can be construed as computing the *root-reachable* subset of control states in a rooted pushdown system, $M = (Q,\Gamma,\delta,q_0)$:

$$\left\{q : q_0 \underset{M}{\longmapsto\!\!\!\!\!\rightarrow} q\right\}.$$

Reps *et. al* and many others provide a straightforward "summarization" algorithm to compute this set (Bouajjani et al. 1997; Kodumal and Aiken 2004; Reps 1998; Reps et al. 2005). We will develop a complete alternative to summarization, and then instrument this development for introspective pushdown systems. Summarization builds two large tables:

- One maps "calling contexts" to "return sites" (AKA "local continuations") so that a returning function steps to all the places it must return to.
- The other maps "calling contexts" to "return states," so that any place performing a call with an already analyzed calling context can jump straight to the returns.

This setup requires intimate knowledge of the language in question for where continuations should be segmented to be "local" and is strongly tied to function call and return. Our algorithm is based on graph traversals of the transition relation for a generic pushdown system. It requires no specialized knowledge of the analyzed language, and it avoids the memory footprint of summary tables.

2.6 Pushdown automata

A **pushdown automaton** is an input-accepting generalization of a rooted pushdown system, a 7-tuple $(Q,\Sigma,\Gamma,\delta,q_0,F,\vec{\gamma})$ in which:

1. Σ is an input alphabet;
2. $\delta \subseteq Q \times \Gamma_{\pm} \times (\Sigma \cup \{\varepsilon\}) \times Q$ is a transition relation;

3. $F \subseteq Q$ is a set of accepting states; and
4. $\vec{\gamma} \in \Gamma^*$ is the initial stack.

We use $\mathbb{PDA}$ to denote the class of all pushdown automata.

Pushdown automata recognize languages over their input alphabet. To do so, their transition relation may optionally consume an input character upon transition. Formally, a PDA $M = (Q, \Sigma, \Gamma, \delta, q_0, F, \vec{\gamma})$ recognizes the language $\mathscr{L}(M) \subseteq \Sigma^*$:

$$\begin{aligned}
\varepsilon &\in \mathscr{L}(M) \text{ if } q_0 \in F \\
aw &\in \mathscr{L}(M) \text{ if } \delta(q_0, \gamma_+, a, q') \text{ and } w \in \mathscr{L}(Q, \Sigma, \Gamma, \delta, q', F, \gamma : \vec{\gamma}) \\
aw &\in \mathscr{L}(M) \text{ if } \delta(q_0, \varepsilon, a, q') \text{ and } w \in \mathscr{L}(Q, \Sigma, \Gamma, \delta, q', F, \vec{\gamma}) \\
aw &\in \mathscr{L}(M) \text{ if } \delta(q_0, \gamma_-, a, q') \text{ and } w \in \mathscr{L}(Q, \Sigma, \Gamma, \delta, q', F, \vec{\gamma}') \\
&\qquad \text{where } \vec{\gamma} = \langle \gamma, \gamma_2, \ldots, \gamma_n \rangle \text{ and } \vec{\gamma}' = \langle \gamma_2, \ldots, \gamma_n \rangle,
\end{aligned}$$

where a is either the empty string ε or a single character.

2.7 Nondeterministic finite automata

In this work, we will need a finite description of all possible stacks at a given control state within a rooted pushdown system. We will exploit the fact that the set of stacks at a given control point is a regular language. Specifically, we will extract a nondeterministic finite automaton accepting that language from the structure of a rooted pushdown system. A **nondeterministic finite automaton** (NFA) is a quintuple $M = (Q, \Sigma, \delta, q_0, F)$:

- Q is a finite set of control states;
- Σ is an input alphabet;
- $\delta \subseteq Q \times (\Sigma \cup \{\varepsilon\}) \times Q$ is a transition relation.
- q_0 is a distinguished start state.
- $F \subseteq Q$ is a set of accepting states.

We denote the class of all NFAs as $\mathbb{NFA}$.

3 Setting: A-Normal Form λ-Calculus

Since our goal is analysis of *higher-order languages*, we operate on the λ-calculus. To simplify presentation of the concrete and abstract semantics, we choose A-Normal Form λ-calculus. (This is a strictly cosmetic choice: all of our results can be replayed *mutatis mutandis* in the standard direct-style setting as well. This differs from CFA2's requirement of CPS, since ANF can be applied locally whereas CPS requires a global transformation.) ANF enforces an order of evaluation and it requires that all arguments to a function be

$$
\begin{array}{rll}
c \in \mathit{Conf} = \mathsf{Exp} \times \mathit{Env} \times \mathit{Store} \times \mathit{Kont} & & \text{[configurations]} \\
\rho \in \mathit{Env} = \mathsf{Var} \rightharpoonup \mathit{Addr} & & \text{[environments]} \\
\sigma \in \mathit{Store} = \mathit{Addr} \to \mathit{Clo} & & \text{[stores]} \\
\mathit{clo} \in \mathit{Clo} = \mathsf{Lam} \times \mathit{Env} & & \text{[closures]} \\
\kappa \in \mathit{Kont} = \mathit{Frame}^* & & \text{[continuations]} \\
\phi \in \mathit{Frame} = \mathsf{Var} \times \mathsf{Exp} \times \mathit{Env} & & \text{[stack frames]} \\
a \in \mathit{Addr} \text{ is an infinite set of addresses} & & \text{[addresses].}
\end{array}
$$

Fig. 3: The concrete configuration-space.

atomic:

$$
\begin{array}{rll}
e \in \mathsf{Exp} ::= & \texttt{(let ((}v\ \mathit{call}\texttt{)) }e\texttt{)} & \text{[non-tail call]} \\
\mid & \mathit{call} & \text{[tail call]} \\
\mid & æ & \text{[return]} \\
f, æ \in \mathsf{Atom} ::= & v \mid \mathit{lam} & \text{[atomic expressions]} \\
\mathit{lam} \in \mathsf{Lam} ::= & \texttt{(}\lambda\ \texttt{(}v\texttt{) }e\texttt{)} & \text{[lambda terms]} \\
\mathit{call} \in \mathsf{Call} ::= & \texttt{(}f\ æ\texttt{)} & \text{[applications]} \\
v \in \mathsf{Var} \text{ is a set of identifiers} & & \text{[variables].}
\end{array}
$$

We use the CESK machine of Felleisen and Friedman (1987) to specify a small-step semantics for ANF. The CESK machine has an explicit stack, and under a structural abstraction, the stack component of this machine directly becomes the stack component of a pushdown system. The set of configurations (*Conf*) for this machine has the four expected components (Figure 3).

3.1 Semantics

To define the semantics, we need five items:

1. $\mathscr{I} : \mathsf{Exp} \to \mathit{Conf}$ injects an expression into a configuration:
$$c_0 = \mathscr{I}(e) = (e, [], [], \langle\rangle).$$
2. $\mathscr{A} : \mathsf{Atom} \times \mathit{Env} \times \mathit{Store} \rightharpoonup \mathit{Clo}$ evaluates atomic expressions:
$$
\begin{array}{rl}
\mathscr{A}(\mathit{lam}, \rho, \sigma) = (\mathit{lam}, \rho) & \text{[closure creation]} \\
\mathscr{A}(v, \rho, \sigma) = \sigma(\rho(v)) & \text{[variable look-up].}
\end{array}
$$
3. $(\Rightarrow) \subseteq \mathit{Conf} \times \mathit{Conf}$ transitions between configurations. (Defined below.)
4. $\mathscr{E} : \mathsf{Exp} \to \mathscr{P}(\mathit{Conf})$ computes the set of reachable machine configurations for a given program:
$$\mathscr{E}(e) = \{c : \mathscr{I}(e) \Rightarrow^* c\}.$$

5. *alloc* : $\mathsf{Var} \times \mathit{Conf} \to \mathit{Addr}$ chooses fresh store addresses for newly bound variables. The address-allocation function is an opaque parameter in this semantics, so that the forthcoming abstract semantics may also parameterize allocation. The nondeterministic nature of the semantics makes any choice of *alloc* sound (Might and Manolios 2009). This parameterization provides the knob to tune the polyvariance and context-sensitivity of the resulting analysis. For the sake of defining the concrete semantics, letting addresses be natural numbers suffices. The allocator can then choose the lowest unused address:

$$\mathit{Addr} = \mathbb{N}$$
$$\mathit{alloc}(v, (e, \rho, \sigma, \kappa)) = 1 + \max(\mathit{dom}(\sigma)).$$

Transition relation To define the transition $c \Rightarrow c'$, we need three rules. The first rule handles tail calls by evaluating the function into a closure, evaluating the argument into a value and then moving to the body of the closure's λ-term:

$$\overbrace{(\llbracket (f\ \text{æ}) \rrbracket, \rho, \sigma, \kappa)}^{c} \Rightarrow \overbrace{(e, \rho'', \sigma', \kappa)}^{c'}, \text{ where}$$
$$(\llbracket (\lambda\ (v)\ e) \rrbracket, \rho') = \mathscr{A}(f, \rho, \sigma)$$
$$a = \mathit{alloc}(v, c)$$
$$\rho'' = \rho'[v \mapsto a]$$
$$\sigma' = \sigma[a \mapsto \mathscr{A}(\text{æ}, \rho, \sigma)].$$

Non-tail calls push a frame onto the stack and evaluate the call:

$$\overbrace{(\llbracket \texttt{(let ((}v\ \mathit{call}\texttt{))}\ e) \rrbracket, \rho, \sigma, \kappa)}^{c} \Rightarrow \overbrace{(\mathit{call}, \rho, \sigma, (v, e, \rho) : \kappa)}^{c'}.$$

Function return pops a stack frame:

$$\overbrace{(\text{æ}, \rho, \sigma, (v, e, \rho') : \kappa)}^{c} \Rightarrow \overbrace{(e, \rho'', \sigma', \kappa)}^{c'}, \text{ where}$$
$$a = \mathit{alloc}(v, c)$$
$$\rho'' = \rho'[v \mapsto a]$$
$$\sigma' = \sigma[a \mapsto \mathscr{A}(\text{æ}, \rho, \sigma)].$$

4 An Infinite-State Abstract Interpretation

Our first step toward a static analysis is an abstract interpretation into an *infinite* state-space. To achieve a pushdown analysis, we simply abstract away less than we normally would. Specifically, we leave the stack height unbounded.

Figure 4 details the abstract configuration-space. To synthesize it, we force addresses to be a finite set, but crucially, we leave the stack untouched. When we compact the set of addresses into a finite set, the machine may run out of addresses to allocate, and when it does, the pigeon-hole principle will force multiple closures to reside at the same address. As a result, to remain sound we change the range of the store to become a power set in the abstract configuration-space. The abstract transition relation has components analogous to those from the concrete semantics:

$$\hat{c} \in \widehat{Conf} = \mathsf{Exp} \times \widehat{Env} \times \widehat{Store} \times \widehat{Kont} \qquad \text{[configurations]}$$
$$\hat{\rho} \in \widehat{Env} = \mathsf{Var} \rightharpoonup \widehat{Addr} \qquad \text{[environments]}$$
$$\hat{\sigma} \in \widehat{Store} = \widehat{Addr} \to \mathscr{P}\left(\widehat{Clo}\right) \qquad \text{[stores]}$$
$$\widehat{clo} \in \widehat{Clo} = \mathsf{Lam} \times \widehat{Env} \qquad \text{[closures]}$$
$$\hat{\kappa} \in \widehat{Kont} = \widehat{Frame}^* \qquad \text{[continuations]}$$
$$\hat{\phi} \in \widehat{Frame} = \mathsf{Var} \times \mathsf{Exp} \times \widehat{Env} \qquad \text{[stack frames]}$$
$$\hat{a} \in \widehat{Addr} \text{ is a } \textit{finite} \text{ set of addresses} \qquad \text{[addresses].}$$

Fig. 4: The abstract configuration-space.

Program injection The abstract injection function $\hat{\mathscr{I}} : \mathsf{Exp} \to \widehat{Conf}$ pairs an expression with an empty environment, an empty store and an empty stack to create the initial abstract configuration:

$$\hat{c}_0 = \hat{\mathscr{I}}(e) = (e, [], [], \langle\rangle).$$

Atomic expression evaluation The abstract atomic expression evaluator, $\hat{\mathscr{A}} : \mathsf{Atom} \times \widehat{Env} \times \widehat{Store} \to \mathscr{P}(\widehat{Clo})$, returns the value of an atomic expression in the context of an environment and a store; it returns a *set* of abstract closures:

$$\hat{\mathscr{A}}(lam, \hat{\rho}, \hat{\sigma}) = \{(lam, \hat{\rho})\} \qquad \text{[closure creation]}$$
$$\hat{\mathscr{A}}(v, \hat{\rho}, \hat{\sigma}) = \hat{\sigma}(\hat{\rho}(v)) \qquad \text{[variable look-up].}$$

Reachable configurations The abstract program evaluator $\hat{\mathscr{E}} : \mathsf{Exp} \to \mathscr{P}(\widehat{Conf})$ returns all of the configurations reachable from the initial configuration:

$$\hat{\mathscr{E}}(e) = \left\{\hat{c} : \hat{\mathscr{I}}(e) \mathrel{\widehat{\Rightarrow}}^* \hat{c}\right\}.$$

Because there are an infinite number of abstract configurations, a naïve implementation of this function may not terminate. Pushdown analysis provides a way of precisely computing this set and both finitely and compactly representing the result.

Transition relation The abstract transition relation $(\widehat{\Rightarrow}) \subseteq \widehat{Conf} \times \widehat{Conf}$ has three rules, one of which has become nondeterministic. A tail call may fork because there could be multiple abstract closures that it is invoking:

$$\overbrace{([\![(f\ ae)]\!], \hat{\rho}, \hat{\sigma}, \hat{\kappa})}^{\hat{c}} \mathrel{\widehat{\Rightarrow}} \overbrace{(e, \hat{\rho}'', \hat{\sigma}', \hat{\kappa})}^{\hat{c}'}, \text{ where}$$
$$([\![(\lambda\ (v)\ e)]\!], \hat{\rho}') \in \hat{\mathscr{A}}(f, \hat{\rho}, \hat{\sigma})$$
$$\hat{a} = \widehat{alloc}(v, \hat{c})$$
$$\hat{\rho}'' = \hat{\rho}'[v \mapsto \hat{a}]$$
$$\hat{\sigma}' = \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathscr{A}}(ae, \hat{\rho}, \hat{\sigma})].$$

We define all of the partial orders shortly, but for stores:

$$(\hat{\sigma} \sqcup \hat{\sigma}')(\hat{a}) = \hat{\sigma}(\hat{a}) \cup \hat{\sigma}'(\hat{a}).$$

A non-tail call pushes a frame onto the stack and evaluates the call:

$$\overbrace{(\llbracket \texttt{(let ((}v\ call\texttt{)) }e\texttt{)} \rrbracket, \hat{\rho}, \hat{\sigma}, \hat{\kappa})}^{\hat{c}} \widehat{\Rightarrow} \overbrace{(call, \hat{\rho}, \hat{\sigma}, (v, e, \hat{\rho}) : \hat{\kappa})}^{\hat{c}'}.$$

A function return pops a stack frame:

$$\overbrace{(æ, \hat{\rho}, \hat{\sigma}, (v, e, \hat{\rho}') : \hat{\kappa})}^{\hat{c}} \widehat{\Rightarrow} \overbrace{(e, \hat{\rho}'', \hat{\sigma}', \hat{\kappa})}^{\hat{c}'}, \text{ where} \qquad \begin{aligned} \hat{a} &= \widehat{alloc}(v, \hat{c}) \\ \hat{\rho}'' &= \hat{\rho}'[v \mapsto \hat{a}] \\ \hat{\sigma}' &= \hat{\sigma} \sqcup [\hat{a} \mapsto \hat{\mathscr{A}}(æ, \hat{\rho}, \hat{\sigma})]. \end{aligned}$$

Allocation: Polyvariance and context-sensitivity In the abstract semantics, the abstract allocation function $\widehat{alloc} : \mathsf{Var} \times \widehat{Conf} \to \widehat{Addr}$ determines the polyvariance of the analysis. In a control-flow analysis, *polyvariance* literally refers to the number of abstract addresses (variants) there are for each variable. An advantage of this framework over CFA2 is that varying this abstract allocation function instantiates pushdown versions of classical flow analyses. All of the following allocation approaches can be used with the abstract semantics. Note, though only a technical detail, that the concrete address space and allocation would change as well for the abstraction function to still work. The abstract allocation function is a parameter to the analysis.

Monovariance: Pushdown 0CFA Pushdown 0CFA uses variables themselves for abstract addresses:

$$\widehat{Addr} = \mathsf{Var}$$
$$alloc(v, \hat{c}) = v.$$

For better precision, a program would be transformed to have unique binders.

Context-sensitive: Pushdown 1CFA Pushdown 1CFA pairs the variable with the current expression to get an abstract address:

$$\widehat{Addr} = \mathsf{Var} \times \mathsf{Exp}$$
$$alloc(v, (e, \hat{\rho}, \hat{\sigma}, \hat{\kappa})) = (v, e).$$

For better precision, expressions are often uniquely labeled so that textually equal expressions at different points in the program are distinguished.

Polymorphic splitting: Pushdown poly/CFA Assuming we compiled the program from a programming language with let expressions and we marked which identifiers were let-bound, we can enable polymorphic splitting:

$$\widehat{Addr} = \mathsf{Var} + \mathsf{Var} \times \mathsf{Exp}$$

$$alloc(v, (\llbracket (f\ \text{æ}) \rrbracket, \hat{\rho}, \hat{\sigma}, \hat{\kappa})) = \begin{cases} (v, \llbracket (f\ \text{æ}) \rrbracket) & f \text{ is let-bound} \\ v & \text{otherwise.} \end{cases}$$

Pushdown k-CFA For pushdown k-CFA, we need to look beyond the current state and at the last k states, necessarily changing the signature of $\widehat{alloc}$ to $\mathsf{Var} \times \widehat{Conf}^* \to \widehat{Addr}$. By concatenating the expressions in the last k states together, and pairing this sequence with a variable we get pushdown k-CFA:

$$\widehat{Addr} = \mathsf{Var} \times \mathsf{Exp}^k$$

$$\widehat{alloc}(v, \langle (e_1, \hat{\rho}_1, \hat{\sigma}_1, \hat{\kappa}_1), \ldots \rangle) = (v, \langle e_1, \ldots, e_k \rangle).$$

4.1 Partial orders

For each set $\hat{X}$ inside the abstract configuration-space, we use the natural partial order, $(\sqsubseteq_{\hat{X}}) \subseteq \hat{X} \times \hat{X}$. Abstract addresses and syntactic sets have flat partial orders. For the other sets, the partial order lifts:

- point-wise over environments:

$$\hat{\rho} \sqsubseteq \hat{\rho}' \text{ iff } \hat{\rho}(v) = \hat{\rho}'(v) \text{ for all } v \in dom(\hat{\rho});$$

- component-wise over closures:

$$(lam, \hat{\rho}) \sqsubseteq (lam, \hat{\rho}') \text{ iff } \hat{\rho} \sqsubseteq \hat{\rho}';$$

- point-wise over stores:

$$\hat{\sigma} \sqsubseteq \hat{\sigma}' \text{ iff } \hat{\sigma}(\hat{a}) \sqsubseteq \hat{\sigma}'(\hat{a}) \text{ for all } \hat{a} \in dom(\hat{\sigma});$$

- component-wise over frames:

$$(v, e, \hat{\rho}) \sqsubseteq (v, e, \hat{\rho}') \text{ iff } \hat{\rho} \sqsubseteq \hat{\rho}';$$

- element-wise over continuations:

$$\langle \hat{\phi}_1, \ldots, \hat{\phi}_n \rangle \sqsubseteq \langle \hat{\phi}'_1, \ldots, \hat{\phi}'_n \rangle \text{ iff } \hat{\phi}_i \sqsubseteq \hat{\phi}'_i; \text{ and}$$

- component-wise across configurations:

$$(e, \hat{\rho}, \hat{\sigma}, \hat{\kappa}) \sqsubseteq (e, \hat{\rho}', \hat{\sigma}', \hat{\kappa}') \text{ iff } \hat{\rho} \sqsubseteq \hat{\rho}' \text{ and } \hat{\sigma} \sqsubseteq \hat{\sigma}' \text{ and } \hat{\kappa} \sqsubseteq \hat{\kappa}'.$$

4.2 Soundness

To prove soundness, an abstraction map α connects the concrete and abstract configuration-spaces:

$$\begin{aligned}
\alpha(e,\rho,\sigma,\kappa) &= (e,\alpha(\rho),\alpha(\sigma),\alpha(\kappa)) \\
\alpha(\rho) &= \lambda v.\alpha(\rho(v)) \\
\alpha(\sigma) &= \lambda \hat{a}. \bigsqcup_{\alpha(a)=\hat{a}} \{\alpha(\sigma(a))\} \\
\alpha\langle\phi_1,\ldots,\phi_n\rangle &= \langle\alpha(\phi_1),\ldots,\alpha(\phi_n)\rangle \\
\alpha(v,e,\rho) &= (v,e,\alpha(\rho)) \\
\alpha(a) &\text{ is determined by the allocation functions.}
\end{aligned}$$

It is then easy to prove that the abstract transition relation simulates the concrete transition relation:

Theorem 4.1
If $\alpha(c) \sqsubseteq \hat{c}$ and $c \Rightarrow c'$, then there exists $\hat{c}' \in \widehat{Conf}$ such that $\alpha(c') \sqsubseteq \hat{c}'$ and $\hat{c} \mathrel{\widehat{\Rightarrow}} \hat{c}'$.

Proof
The proof follows by case analysis on the expression in the configuration. It is a straightforward adaptation of similar proofs, such as that of Might (2007) for k-CFA. □

5 From the Abstracted CESK Machine to a PDA

In the previous section, we constructed an infinite-state abstract interpretation of the CESK machine. The infinite-state nature of the abstraction makes it difficult to see how to answer static analysis questions. Consider, for instance, a control flow-question:

At the call site $(f\ æ)$*, may a closure over lam be called?*

If the abstracted CESK machine were a finite-state machine, an algorithm could answer this question by enumerating all reachable configurations and looking for an abstract configuration $([\![(f\ æ)]\!],\hat{\rho},\hat{\sigma},\hat{\kappa})$ in which $(lam,_) \in \hat{\mathscr{A}}(f,\hat{\rho},\hat{\sigma})$. However, because the abstracted CESK machine may contain an infinite number of reachable configurations, an algorithm cannot enumerate them.

Fortunately, we can recast the abstracted CESK as a special kind of infinite-state system: a pushdown automaton (PDA). Pushdown automata occupy a sweet spot in the theory of computation: they have an infinite configuration-space, yet many useful properties (*e.g.*, word membership, non-emptiness, control-state reachability) remain decidable. Once the abstracted CESK machine becomes a PDA, we can answer the control-flow question by checking whether a specific regular language, accounting for the states of interest, when intersected with the language of the PDA, is nonempty.

The recasting as a PDA is a shift in perspective. A configuration has an expression, an environment and a store. A stack character is a frame. We choose to make the alphabet the set of control states, so that the language accepted by the PDA will be sequences of control-states visited by the abstracted CESK machine. Thus, every transition will consume the

$$\widehat{\mathscr{PDA}}(e) = (Q, \Sigma, \Gamma, \delta, q_0, F, \langle\rangle), \text{ where}$$
$$Q = \mathsf{Exp} \times \widehat{Env} \times \widehat{Store}$$
$$\Sigma = Q$$
$$\Gamma = \widehat{Frame}$$

$$(q, \varepsilon, q', q') \in \delta \text{ iff } (q, \hat{\kappa}) \widehat{\Rightarrow} (q', \hat{\kappa}) \text{ for all } \hat{\kappa}$$
$$(q, \hat{\phi}_-, q', q') \in \delta \text{ iff } (q, \hat{\phi} : \hat{\kappa}) \widehat{\Rightarrow} (q', \hat{\kappa}) \text{ for all } \hat{\kappa}$$
$$(q, \hat{\phi}'_+, q', q') \in \delta \text{ iff } (q, \hat{\kappa}) \widehat{\Rightarrow} (q', \hat{\phi}' : \hat{\kappa}) \text{ for all } \hat{\kappa}$$
$$(q_0, \langle\rangle) = \hat{\mathscr{I}}(e)$$
$$F = Q.$$

Fig. 5: $\widehat{\mathscr{PDA}} : \mathsf{Exp} \to \mathbb{PDA}$.

control-state to which it transitioned as an input character. Figure 5 defines the program-to-PDA conversion function $\widehat{\mathscr{PDA}} : \mathsf{Exp} \to \mathbb{PDA}$. (Note the implicit use of the isomorphism $Q \times \widehat{Kont} \cong \widehat{Conf}$.)

At this point, we can answer questions about whether a specified control state is reachable by formulating a question about the intersection of a regular language with a context-free language described by the PDA. That is, if we want to know whether the control state $(e', \hat{\rho}, \hat{\sigma})$ is reachable in a program e, we can reduce the problem to determining:

$$\Sigma^* \cdot \{(e', \hat{\rho}, \hat{\sigma})\} \cdot \Sigma^* \cap \mathscr{L}(\widehat{\mathscr{PDA}}(e)) \neq \emptyset,$$

where $L_1 \cdot L_2$ is the concatenation of formal languages L_1 and L_2.

Theorem 5.1
Control-state reachability is decidable.

Proof
The intersection of a regular language and a context-free language is context-free (simple machine product of PDA with DFA). The emptiness of a context-free language is decidable. The decision procedure is easiest for CFGs: mark terminals, mark non-terminals that reduce to marked (non)terminals until we reach a fixed point. If the start symbol is marked, then the language is nonempty. The PDA to CFG translation is a standard construction. □

Now, consider how to use control-state reachability to answer the control-flow question from earlier. There are a finite number of possible control states in which the λ-term *lam* may flow to the function f in call site $(f\ æ)$; let's call this set of states $\hat{S}$:

$$\hat{S} = \{([\![(f\ æ)]\!], \hat{\rho}, \hat{\sigma}) : (lam, \hat{\rho}') \in \hat{\mathscr{A}}(f, \hat{\rho}, \hat{\sigma}) \text{ for some } \hat{\rho}'\}.$$

What we want to know is whether any state in the set $\hat{S}$ is reachable in the PDA. In effect what we are asking is whether there exists a control state $q \in \hat{S}$ such that:

$$\Sigma^* \cdot \{q\} \cdot \Sigma^* \cap \mathscr{L}(\widehat{\mathscr{PDA}}(e)) \neq \emptyset.$$

If this is true, then *lam* may flow to f; if false, then it does not.

Problem: Doubly exponential complexity The non-emptiness-of-intersection approach establishes decidability of pushdown control-flow analysis. But, two exponential complexity barriers make this technique impractical.

First, there are an exponential number of both environments ($|\widehat{Addr}|^{|\mathsf{Var}|}$) and stores ($(2^{|\widehat{Clo}|})^{|\widehat{Addr}|} = 2^{|\widehat{Clo}|\times|\widehat{Addr}|}$) to consider for the set $\hat{S}$. On top of that, computing the intersection of a regular language with a context-free language will require enumeration of the (exponential) control-state-space of the PDA. The size of the control-state-space of the PDA is clearly doubly exponential:

$$\begin{aligned}
|Q| &= |\mathsf{Exp} \times \widehat{Env} \times \widehat{Store}| \\
&= |\mathsf{Exp}| \times |\widehat{Env}| \times |\widehat{Store}| \\
&= |\mathsf{Exp}| \times |\widehat{Addr}|^{|\mathsf{Var}|} \times 2^{|\widehat{Clo}|\times|\widehat{Addr}|} \\
&= |\mathsf{Exp}| \times |\widehat{Addr}|^{|\mathsf{Var}|} \times 2^{|\mathsf{Lam}\times\widehat{Env}|\times|\widehat{Addr}|} \\
&= |\mathsf{Exp}| \times |\widehat{Addr}|^{|\mathsf{Var}|} \times 2^{|\mathsf{Lam}|\times|\widehat{Addr}|^{|\mathsf{Var}|}\times|\widehat{Addr}|}
\end{aligned}$$

As a result, this approach is doubly exponential. For the next few sections, our goal will be to lower the complexity of pushdown control-flow analysis.

6 Focusing on Reachability

In the previous section, we saw that control-flow analysis reduces to the reachability of certain control states within a pushdown system. We also determined reachability by converting the abstracted CESK machine into a PDA, and using emptiness-testing on a language derived from that PDA. Unfortunately, we also found that this approach is deeply exponential.

Since control-flow analysis reduced to the reachability of control-states in the PDA, we skip the language problems and go directly to reachability algorithms of Bouajjani et al. (1997); Kodumal and Aiken (2004); Reps (1998) and Reps et al. (2005) that determine the reachable *configurations* within a pushdown system. These algorithms are even polynomial-time. Unfortunately, some of them are polynomial-time in the number of control states, and in the abstracted CESK machine, there are an exponential number of control states. We don't want to *enumerate* the entire control state-space, or else the search becomes exponential in even the best case.

To avoid this worst-case behavior, we present a straightforward pushdown-reachability algorithm that considers only the *reachable* control states. We cast our reachability algorithm as a fixed-point iteration, in which we incrementally construct the reachable subset of a pushdown system. A rooted pushdown system $M = (Q, \Gamma, \delta, q_0)$ is *compact* if for any $(q, g, q') \in \delta$, it is the case that:

$$(q_0, \langle\rangle) \underset{M}{\longmapsto}^* (q, \vec{\gamma}) \text{ for some stack } \vec{\gamma},$$

and the domain of states and stack characters are exactly those that appear in δ:

$$\begin{aligned}
Q &= \bigcup \{\{q, q'\} : (q, g, q') \in \delta\} \\
\Gamma &= \{\gamma : (q, \gamma_+, q') \in \delta \text{ or } (q, \gamma_-, q') \in \delta\}
\end{aligned}$$

In other words, a rooted pushdown system is compact when its states, transitions and stack characters appear on legal paths from the initial control state. We will refer to the class of compact rooted pushdown systems as $\mathbb{CRPDS}$.

We can compact a rooted pushdown system with a map:

$$\mathscr{C} : \mathbb{RPDS} \to \mathbb{CRPDS}$$

$$\mathscr{C}\overbrace{(Q,\Gamma,\delta,q_0)}^{M} = (Q',\Gamma',\delta',q_0)$$

$$\text{where } Q' = \left\{ q : (q_0,\langle\rangle) \underset{M}{\longmapsto}^* (q,\vec{\gamma}) \right\}$$

$$\Gamma' = \left\{ \gamma_i : (q_0,\langle\rangle) \underset{M}{\longmapsto}^* (q,\vec{\gamma}) \right\}$$

$$\delta' = \left\{ (q,g,q') : (q_0,\langle\rangle) \underset{M}{\overset{\vec{g}}{\longmapsto}} (q,[\vec{g}]) \underset{M}{\overset{g}{\longmapsto}} (q',[\vec{g}g]) \right\}.$$

In practice, the real difference between a rooted pushdown system and its compact form is that our original system will be defined intensionally (having come from the components of an abstracted CESK machine), whereas the compact system will be defined extensionally, with the contents of each component explicitly enumerated during its construction.

Our near-term goals are (1) to convert our abstracted CESK machine into a rooted pushdown system and (2) to find an *efficient* method to compact it.

To convert the abstracted CESK machine into a rooted pushdown system, we use the function $\widehat{\mathscr{RPDS}} : \mathsf{Exp} \to \mathbb{RPDS}$:

$$\widehat{\mathscr{RPDS}}(e) = (Q,\Gamma,\delta,q_0)$$

$$Q = \mathsf{Exp} \times \widehat{Env} \times \widehat{Store}$$

$$\Gamma = \widehat{Frame}$$

$$(q_0,\langle\rangle) = \hat{\mathscr{I}}(e)$$

$$q \overset{\varepsilon}{\rightarrowtail} q' \in \delta \text{ iff } (q,\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (q',\hat{\kappa}) \text{ for all } \hat{\kappa}$$

$$q \overset{\hat{\phi}_-}{\rightarrowtail} q' \in \delta \text{ iff } (q,\hat{\phi} : \hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (q',\hat{\kappa}) \text{ for all } \hat{\kappa}$$

$$q \overset{\hat{\phi}_+}{\rightarrowtail} q' \in \delta \text{ iff } (q,\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (q',\hat{\phi} : \hat{\kappa}) \text{ for all } \hat{\kappa}.$$

7 Compacting a Rooted Pushdown System

We now turn our attention to compacting a rooted pushdown system (defined intensionally) into its compact form (defined extensionally). That is, we want to find an implementation of the function $\mathscr{C}$. To do so, we first phrase the construction as the least fixed point of a monotonic function. This will provide a method (albeit an inefficient one) for computing the function $\mathscr{C}$. In the next section, we look at an optimized work-set driven algorithm that avoids the inefficiencies of this section's algorithm.

The function $\mathscr{F} : \mathbb{RPDS} \to (\mathbb{CRPDS} \to \mathbb{CRPDS})$ generates the monotonic iteration function we need:

$$\mathscr{F}(M) = f, \text{ where } M = (Q,\Gamma,\delta,q_0)$$

$$f(S,\Gamma,E,s_0) = (S',\Gamma,E',s_0), \text{ where}$$

$$S' = S \cup \left\{ s' : s \in S \text{ and } s \underset{M}{\longmapsto\!\!\!\!\!\rightarrow} s' \right\} \cup \{s_0\}$$

$$E' = E \cup \left\{ s \overset{g}{\rightarrowtail} s' : s \in S \text{ and } s \underset{M}{\overset{g}{\longmapsto\!\!\!\!\!\rightarrow}} s' \right\}.$$

Given a rooted pushdown system M, each application of the function $\mathscr{F}(M)$ accretes new edges at the frontier of the system. Once the algorithm reaches a fixed point, the system is complete:

Theorem 7.1

$\mathscr{C}(M) = \text{lfp}(\mathscr{F}(M))$.

Proof

Let $M = (Q, \Gamma, \delta, q_0)$. Let $f = \mathscr{F}(M)$. Observe that $\text{lfp}(f) = f^n(\emptyset, \Gamma, \emptyset, q_0)$ for some n. When $N \subseteq \mathscr{C}(M)$, then it easy to show that $f(N) \subseteq \mathscr{C}(M)$. Hence, $\mathscr{C}(M) \supseteq \text{lfp}(\mathscr{F}(M))$.

To show $\mathscr{C}(M) \subseteq \text{lfp}(\mathscr{F}(M))$, suppose this is not the case. Then, there must be at least one edge in $\mathscr{C}(M)$ that is not in $\text{lfp}(\mathscr{F}(M))$. Since these edges must be root reachable, let (s, g, s') be the first such edge in some path from the root. This means that the state s is in $\text{lfp}(\mathscr{F}(M))$. Let m be the lowest natural number such that s appears in $f^m(M)$. By the definition of f, this edge must appear in $f^{m+1}(M)$, which means it must also appear in $\text{lfp}(\mathscr{F}(M))$, which is a contradiction. Hence, $\mathscr{C}(M) \subseteq \text{lfp}(\mathscr{F}(M))$. □

7.1 Complexity: Polynomial and exponential

To determine the complexity of this algorithm, we ask two questions: how many times would the algorithm invoke the iteration function in the worst case, and how much does each invocation cost in the worst case? The size of the final system bounds the run-time of the algorithm. Suppose the final system has m states. In the worst case, the iteration function adds only a single edge each time. Since there are at most $2|\Gamma|m^2 + m^2$ edges in the final graph, the maximum number of iterations is $2|\Gamma|m^2 + m^2$.

The cost of computing each iteration is harder to bound. The cost of determining whether to add a push edge is proportional to the size of the stack alphabet, while the cost of determining whether to add an ε-edge is constant, so the cost of determining all new push and ε edges to add is proportional to $|\Gamma|m + m$. Determining whether or not to add a pop edge is expensive. To add the pop edge $s \overset{\gamma_-}{\rightarrowtail} s'$, we must prove that there exists a configuration-path to the control state s, in which the character γ is on the top of the stack. This reduces to a CFL-reachability query (Melski and Reps 2000) at each node, the cost of which is $O(|\Gamma_\pm|^3 m^3)$ (Kodumal and Aiken 2004).

To summarize, in terms of the number of reachable control states, the complexity of the most recent algorithm is:

$$O((2|\Gamma|m^2 + m^2) \times (|\Gamma|m + m + |\Gamma_\pm|^3 m^3)) = O(|\Gamma|^4 m^5).$$

While this approach is polynomial in the number of reachable control states, it is far from efficient. In the next section, we provide an optimized version of this fixed-point algorithm that maintains a work-set and an ε-closure graph to avoid spurious recomputation.

Moreover, we have carefully phrased the complexity in terms of "reachable" control states because, in practice, compact rooted pushdown systems will be extremely sparse, and because the maximum number of control states is exponential in the size of the input program. After the subsequent refinement, we will be able to develop a hierarchy of push-down control-flow analyses that employs widening to achieve a polynomial-time algorithm at its foundation.

8 An Efficient Algorithm: Work-sets and ε-Closure Graphs

We have developed a fixed-point formulation of the rooted pushdown system compaction algorithm, but found that, in each iteration, it wasted effort by passing over all discovered states and edges, even though most will not contribute new states or edges. Taking a cue from graph search, we can adapt the fixed-point algorithm with a work-set. That is, our next algorithm will keep a work-set of new states and edges to consider, instead of reconsidering all of them. We will refer to the compact rooted pushdown system we are constructing as a graph, since that is how we represent it (Q is the set of nodes, and δ is a set of labeled edges).

In each iteration, it will pull new states and edges from the work list, insert them into the graph and then populate the work-set with new states and edges that have to be added as a consequence of the recent additions.

8.1 ε-closure graphs

Figuring out what edges to add as a consequence of another edge requires care, for adding an edge can have ramifications on distant control states. Consider, for example, adding the ε-edge $q \rightarrowtail^{\varepsilon} q'$ into the following graph:

$$q_0 \xrightarrow{\gamma_+} q \qquad q' \xrightarrow{\gamma_-} q_1$$

As soon this edge drops in, an ε-edge "implicitly" appears between q_0 and q_1 because the net stack change between them is empty; the resulting graph looks like:

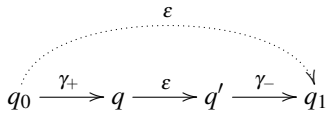

where we have illustrated the implicit ε-edge as a dotted line.

To keep track of these implicit edges, we will construct a second graph in conjunction with the graph: an ε-closure graph. In the ε-closure graph, every edge indicates the existence of a no-net-stack-change path between control states. The ε-closure graph simplifies the task of figuring out which states and edges are impacted by the addition of a new edge.

Formally, an **ε-closure graph**, $H \subseteq N \times N$, is a set of edges. Of course, all ε-closure graphs are reflexive: every node has a self loop. We use the symbol $\mathbb{ECG}$ to denote the class of all ε-closure graphs.

We have two notations for finding ancestors and descendants of a state in an ε-closure graph:

$$\overleftarrow{G}_{\varepsilon}[s] = \{s' : (s', s) \in H\} \cup \{s\} \qquad \text{[ancestors]}$$
$$\overrightarrow{G}_{\varepsilon}[s] = \{s' : (s, s') \in H\} \cup \{s\} \qquad \text{[descendants].}$$

8.2 Integrating a work-set

Since we only want to consider new states and edges in each iteration, we need a work-set, or in this case, three work-sets:

- ΔS contains states to add,
- ΔE contains edges to add,
- ΔH contains new ε-edges.

Let $\mathbb{WS} ::= (\Delta S, \Delta E, \Delta H)$ be the space of work-sets.

8.3 A new fixed-point iteration-space

Instead of consuming a graph and producing a graph, the new fixed-point iteration function will consume and produce a graph, an ε-closure graph, and the work-sets. Hence, the iteration space of the new algorithm is:

$$ICRPDS = (\wp(Q) \times \wp(Q \times \Gamma_{\pm} \times Q)) \times \mathbb{ECG} \times \mathbb{WS}.$$

The *I* in *ICRPDS* stands for *intermediate*.

8.4 The ε-closure graph work-list algorithm

The function $\mathscr{F}' : \mathbb{RPDS} \to (ICRPDS \to ICRPDS)$ generates the required iteration function (Figure 6). Please note that we implicitly distribute union across tuples:

$$(X,Y) \cup (X',Y') = (X \cup X, Y \cup Y').$$

The functions *sprout*, *addPush*, *addPop*, *addEmpty* (defined shortly) calculate the additional the graph edges and ε-closure graph edges (potentially) introduced by a new state or edge.

Sprouting Whenever a new state gets added to the graph, the algorithm must check whether that state has any new edges to contribute. Both push edges and ε-edges do not depend on the current stack, so any such edges for a state in the pushdown system's transition function belong in the graph. The sprout function:

$$sprout_{(Q,\Gamma,\delta,s_0)} : Q \to (\mathscr{P}(\delta) \times \mathscr{P}(Q \times Q)),$$

checks whether a new state could produce any new push edges or no-change edges. We can represent its behavior diagrammatically (as previously, the dotted arrows correspond to the corresponding additions to the work-graph and ε-closure work graph):

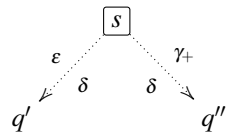

which means if adding control state s:

add edge $s \overset{\varepsilon}{\rightarrowtail} q'$ if it exists in δ (hence the arrow subscript δ), and
add edge $s \overset{\gamma_+}{\rightarrowtail} q''$ if it exists in δ.

Formally:

$$sprout_{(Q,\Gamma,\delta,s_0)}(s) = (\Delta E, \Delta H), \text{ where}$$

$$
\begin{aligned}
\mathscr{F}'(M) &= f, \text{ where} \\
M &= (Q, \Gamma, \delta, q_0) \\
f(G, H, (\Delta S, \Delta E, \Delta H)) &= (G', H', (\Delta S' - S', \Delta E' - E', \Delta H' - H)), \text{ where} \\
(S, \Gamma, E, s_0) &= G \\
(\Delta E_0, \Delta H_0) &= \bigcup_{s \in \Delta S} \mathit{sprout}_M(s) \\
(\Delta E_1, \Delta H_1) &= \bigcup_{(s, \gamma_+, s') \in \Delta E} \mathit{addPush}_M(G, H)(s, \gamma_+, s') \\
(\Delta E_2, \Delta H_2) &= \bigcup_{(s, \gamma_-, s') \in \Delta E} \mathit{addPop}_M(G, H)(s, \gamma_-, s') \\
(\Delta E_3, \Delta H_3) &= \bigcup_{(s, \varepsilon, s') \in \Delta E} \mathit{addEmpty}_M(G, H)(s, s') \\
(\Delta E_4, \Delta H_4) &= \bigcup_{(s, s') \in \Delta H} \mathit{addEmpty}_M(G, H)(s, s') \\
S' &= S \cup \Delta S \\
E' &= E \cup \Delta E \\
H' &= H \cup \Delta H \\
\Delta E' &= \Delta E_0 \cup \Delta E_1 \cup \Delta E_2 \cup \Delta E_3 \cup \Delta E_4 \\
\Delta S' &= \left\{ s' : (s, g, s') \in \Delta E' \right\} \cup \{s_0\} \\
\Delta H' &= \Delta H_0 \cup \Delta H_1 \cup \Delta H_2 \cup \Delta H_3 \cup \Delta H_4 \\
\Delta \Gamma &= \left\{ \gamma : (s, \gamma_+, s') \in \Delta E' \right\} \\
G' &= (S \cup \Delta S, \Gamma \cup \Delta \Gamma, E', q_0).
\end{aligned}
$$

Fig. 6: The fixed point of the function $\mathscr{F}'(M)$ contains the compact form of the rooted pushdown system M.

$$
\begin{aligned}
\Delta E &= \left\{ s \overset{\varepsilon}{\rightarrowtail} q : s \overset{\varepsilon}{\rightarrowtail} q \in \delta \right\} \cup \left\{ s \overset{\gamma_+}{\rightarrowtail} q : s \overset{\gamma_+}{\rightarrowtail} q \in \delta \right\} \\
\Delta H &= \left\{ s \rightarrowtail q : s \overset{\varepsilon}{\rightarrowtail} q \in \delta \right\}.
\end{aligned}
$$

Considering the consequences of a new push edge Once our algorithm adds a new push edge to a graph, there is a chance that it will enable new pop edges for the same stack frame somewhere downstream. If and when it does enable pops, it will also add new edges to the ε-closure graph. The *addPush* function:

$$
\mathit{addPush}_{(Q, \Gamma, \delta, s_0)} : \mathbb{RPDS} \times \mathbb{ECG} \to \delta \to (\mathscr{P}(\delta) \times \mathscr{P}(Q \times Q)),
$$

checks for ε-reachable states that could produce a pop. We can represent this action by the following diagram (the arrow subscript ε indicates edges in the ε-closure graph):

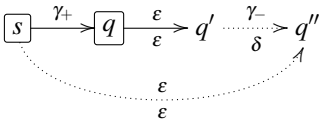

which means if adding push-edge $s \overset{\gamma_+}{\rightarrowtail} q$:

if pop-edge $q' \overset{\gamma_-}{\rightarrowtail} q''$ is in δ, then
add edge $q' \overset{\gamma_-}{\rightarrowtail} q''$, and
add ε-edge $s \rightarrowtail q''$.

Formally:

$$addPush_{(Q,\Gamma,\delta,s_0)}(G,H)(s \overset{\gamma_+}{\rightarrowtail} q) = (\Delta E, \Delta H), \text{ where}$$
$$\Delta E = \left\{ q' \overset{\gamma_-}{\rightarrowtail} q'' : q' \in \overrightarrow{G}_\varepsilon[q] \text{ and } q' \overset{\gamma_-}{\rightarrowtail} q'' \in \delta \right\}$$
$$\Delta H = \left\{ s \rightarrowtail q'' : q' \in \overrightarrow{G}_\varepsilon[q] \text{ and } q' \overset{\gamma_-}{\rightarrowtail} q'' \in \delta \right\}.$$

Considering the consequences of a new pop edge Once the algorithm adds a new pop edge to a graph, it will create at least one new ε-closure graph edge and possibly more by matching up with upstream pushes. The *addPop* function:

$$addPop_{(Q,\Gamma,\delta,s_0)} : \mathbb{RPDS} \times \mathbb{ECG} \to \delta \to (\mathscr{P}(\delta) \times \mathscr{P}(Q \times Q)),$$

checks for ε-reachable push-edges that could match this pop-edge. This action is illustrated by the following diagram:

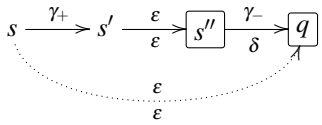

which means if adding pop-edge $s'' \overset{\gamma_-}{\rightarrowtail} q$:

if push-edge $s \overset{\gamma_+}{\rightarrowtail} s'$ is already in the graph, then
add ε-edge $s \rightarrowtail q$.

Formally:

$$addPop_{(Q,\Gamma,\delta,s_0)}(G,H)(s'' \overset{\gamma_-}{\rightarrowtail} q) = (\Delta E, \Delta H), \text{ where}$$
$$\Delta E = \emptyset \text{ and } \Delta H = \left\{ s \rightarrowtail q : s' \in \overleftarrow{G}_\varepsilon[s''] \text{ and } s \overset{\gamma_+}{\rightarrowtail} s' \in G \right\}.$$

Considering the consequences of a new ε-edge Once the algorithm adds a new ε-closure graph edge, it may transitively have to add more ε-closure graph edges, and it may connect

an old push to (perhaps newly enabled) pop edges. The *addEmpty* function:

$$addEmpty_{(Q,\Gamma,\delta,s_0)} : \mathbb{RPDS} \times \mathbb{ECG} \to (Q \times Q) \to (\mathscr{P}(\delta) \times \mathscr{P}(Q \times Q)),$$

checks for newly enabled pops and ε-closure graph edges: Once again, we can represent this action diagrammatically:

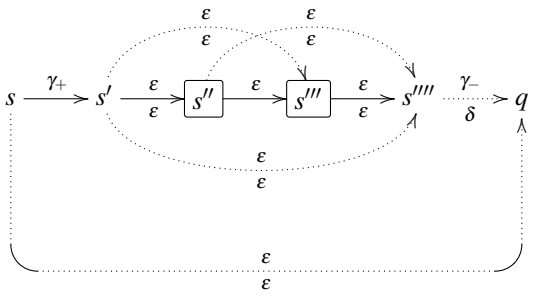

which means if adding ε-edge $s'' \rightarrowtail s'''$:

if pop-edge $s'''' \overset{\gamma_-}{\rightarrowtail} q$ is in δ, then
  add ε-edge $s \rightarrowtail q$; and
  add edge $s'''' \overset{\gamma_-}{\rightarrowtail} q$;
add ε-edges $s' \rightarrowtail s'''$, $s'' \rightarrowtail s''''$, and $s' \rightarrowtail s''''$.

Formally:

$$\begin{aligned}
&addEmpty_{(Q,\Gamma,\delta,s_0)}(G,H)(s'' \rightarrowtail s''') = (\Delta E, \Delta H), \text{ where} \\
&\Delta E = \left\{ s'''' \overset{\gamma_-}{\rightarrowtail} q : s' \in \overleftarrow{G}_\varepsilon[s''] \text{ and } s'''' \in \overrightarrow{G}_\varepsilon[s'''] \text{ and } s \overset{\gamma_+}{\rightarrowtail} s' \in G \text{ and } s'''' \overset{\gamma_-}{\rightarrowtail} q \in \delta \right\} \\
&\Delta H = \left\{ s \rightarrowtail q : s' \in \overleftarrow{G}_\varepsilon[s''] \text{ and } s'''' \in \overrightarrow{G}_\varepsilon[s'''] \text{ and } s \overset{\gamma_+}{\rightarrowtail} s' \in G \text{ and } s'''' \overset{\gamma_-}{\rightarrowtail} q \in \delta \right\} \\
&\quad \cup \left\{ s' \rightarrowtail s''' : s' \in \overleftarrow{G}_\varepsilon[s''] \right\} \\
&\quad \cup \left\{ s'' \rightarrowtail s'''' : s'''' \in \overrightarrow{G}_\varepsilon[s'''] \right\} \\
&\quad \cup \left\{ s' \rightarrowtail s'''' : s' \in \overleftarrow{G}_\varepsilon[s''] \text{ and } s'''' \in \overrightarrow{G}_\varepsilon[s'''] \right\}.
\end{aligned}$$

8.5 Termination and correctness

To prove that a fixed point exists, we show the iteration function is monotonic. The key observation is that ΔG and ΔH drive all additions to, and are disjoint from, G and H. Since G and H monotonically increase in a finite space, ΔG and ΔH run out of room (full details in 18.1). Once the graph reaches a fixed point, all work-sets will be empty, and the ε-closure graph will also be saturated. We can also show that this algorithm is correct by defining first $\mathscr{ECG} : \mathbb{RPDS} \to \mathbb{ECG}$ as

$$\mathscr{ECG}(M) = \left\{ s \rightarrowtail s' : s \overset{\vec{g}}{\underset{M}{\twoheadrightarrow}} s' \text{ and } [\vec{g}] = \varepsilon \right\}$$

and stating the following theorem:

Theorem 8.1
For all $M \in \mathbb{RPDS}$, $\mathscr{C}(M) = G$ and $\mathscr{ECG}(M) = H$, where $(G,H,(\emptyset,\emptyset,\emptyset)) = \mathrm{lfp}(\mathscr{F}'(M))$.

In the proof of Theorem 8.1, the $\subseteq$ case comes from an invariant lemma we have on $\mathscr{F}'$:

Lemma 8.1

$$inv((S,E),H,(\Delta S,\Delta E,\Delta H)) = \forall s \overset{g}{\rightarrowtail} s' \in E \cup \Delta E . s \overset{g}{\underset{M}{\longmapsto\!\!\!\!\!\rightarrow}} s' \\ \wedge \forall s \rightarrowtail s' \in H \cup \Delta H . \exists \vec{g} . [\vec{g}] = \varepsilon \wedge \forall \hat{\kappa} . (s,\hat{\kappa}) \overset{\vec{g}}{\underset{M}{\longmapsto}}{}^{*} (s,\hat{\kappa})$$

The $\supseteq$ case follows from

Lemma 8.2
For all traces $\pi \equiv s_0 \overset{\vec{g}}{\underset{M}{\longmapsto\!\!\!\!\!\rightarrow}} s$, there is both a corresponding path $s_0 \overset{\vec{g}}{\underset{G}{\longmapsto\!\!\!\!\!\rightarrow}} s$ and for all non-empty subtraces of π, $s_b \overset{\vec{g'}}{\underset{M}{\longmapsto\!\!\!\!\!\rightarrow}} s_a$, if $[\vec{g'}] = \varepsilon$ then $s_b \rightarrowtail s_a \in H$.

Since all edges in a compact rooted pushdown system must be in a path from the initial state, we can extract the edges from said paths using this lemma.

8.6 Complexity: Still exponential, but more efficient

As in the previous case (Section 7.1), to determine the complexity of this algorithm, we ask two questions: how many times would the algorithm invoke the iteration function in the worst case, and how much does each invocation cost in the worst case? The run-time of the algorithm is bounded by the size of the final graph plus the size of the ε-closure graph. Suppose the final graph has m states. In the worst case, the iteration function adds only a single edge each time. There are at most $2|\Gamma|m^2 + m^2$ edges in the graph ($|\Gamma|m^2$ push edges, just as many pop edges, and m^2 no-change edges) and at most m^2 edges in the ε-closure graph, which bounds the number of iterations. Recall that m can be exponential in the size of the program, since $m \leq |Q|$ (and Section 5 derived the exponential size of $|Q|$).

Next, we must reason about the worst-case cost of adding an edge: how many edges might an individual iteration consider? In the worst case, the algorithm will consider every edge in every iteration, leading to an asymptotic time-complexity of:

$$O((2|\Gamma|m^2 + 2m^2)^2) = O(|\Gamma|^2 m^4).$$

While still high, this is a an improvement upon the previous algorithm. For sparse graphs, this is a reasonable algorithm.

9 Polynomial-Time Complexity from Widening

In the previous section, we developed a more efficient fixed-point algorithm for computing a compact rooted pushdown system. Even with the core improvements we made, the algorithm remained exponential in the worst case, owing to the fact that there could

be an exponential number of reachable control states. When an abstract interpretation is intolerably complex, the standard approach for reducing complexity and accelerating convergence is widening (Cousot and Cousot 1977). Of course, widening techniques trade away some precision to gain this speed. It turns out that the small-step variants of finite-state CFAs are exponential without some sort of widening as well (Van Horn and Mairson 2008).

To achieve polynomial time complexity for pushdown control-flow analysis requires the same two steps as the classical case: (1) widening the abstract interpretation to use a global, "single-threaded" store and (2) selecting a monovariant allocation function to collapse the abstract configuration-space. Widening eliminates a source of exponentiality in the size of the store; monovariance eliminates a source of exponentiality from environments. In this section, we redevelop the pushdown control-flow analysis framework with a single-threaded store and calculate its complexity.

9.1 Step 1: Refactor the concrete semantics

First, consider defining the reachable states of the concrete semantics using fixed points. That is, let the system-space of the evaluation function be sets of configurations:

$$C \in \mathit{System} = \mathscr{P}(\mathit{Conf}) = \mathscr{P}(\mathsf{Exp} \times \mathit{Env} \times \mathit{Store} \times \mathit{Kont}).$$

We can redefine the concrete evaluation function:

$$\begin{aligned}&\mathscr{E}(e) = \mathrm{lfp}(f_e), \text{ where } f_e : \mathit{System} \to \mathit{System} \text{ and}\\ &f_e(C) = \{\mathscr{I}(e)\} \cup \left\{c' : c \in C \text{ and } c \Rightarrow c'\right\}.\end{aligned}$$

9.2 Step 2: Refactor the abstract semantics

We can take the same approach with the abstract evaluation function, first redefining the abstract system-space:

$$\begin{aligned}\hat{C} \in \widehat{\mathit{System}} &= \mathscr{P}\left(\widehat{\mathit{Conf}}\right)\\ &= \mathscr{P}\left(\mathsf{Exp} \times \widehat{\mathit{Env}} \times \widehat{\mathit{Store}} \times \widehat{\mathit{Kont}}\right),\end{aligned}$$

and then the abstract evaluation function:

$$\begin{aligned}&\hat{\mathscr{E}}(e) = \mathrm{lfp}(\hat{f}_e), \text{ where } \hat{f}_e : \widehat{\mathit{System}} \to \widehat{\mathit{System}} \text{ and}\\ &\hat{f}_e(\hat{C}) = \left\{\hat{\mathscr{I}}(e)\right\} \cup \left\{\hat{c}' : \hat{c} \in \hat{C} \text{ and } \hat{c} \mathrel{\widehat{\Rightarrow}} \hat{c}'\right\}.\end{aligned}$$

What we'd like to do is shrink the abstract system-space with a refactoring that corresponds to a widening.

9.3 Step 3: Single-thread the abstract store

We can approximate a set of abstract stores $\{\hat{\sigma}_1, \ldots, \hat{\sigma}_n\}$ with the least-upper-bound of those stores: $\hat{\sigma}_1 \sqcup \cdots \sqcup \hat{\sigma}_n$. We can exploit this by creating a new abstract system space in which the store is factored out of every configuration. Thus, the system-space contains a

set of *partial configurations* and a single global store:

$$\widehat{System}' = \mathscr{P}\left(\widehat{PConf}\right) \times \widehat{Store}$$
$$\hat{\pi} \in \widehat{PConf} = \mathsf{Exp} \times \widehat{Env} \times \widehat{Kont}.$$

We can factor the store out of the abstract transition relation as well, so that $(\rightarrowtriangle^{\hat{\sigma}}) \subseteq \widehat{PConf} \times (\widehat{PConf} \times \widehat{Store})$:

$$(e, \hat{\rho}, \hat{\kappa}) \overset{\hat{\sigma}}{\rightarrowtriangle} ((e', \hat{\rho}', \hat{\kappa}'), \hat{\sigma}') \text{ iff } (e, \hat{\rho}, \hat{\sigma}, \hat{\kappa}) \widehat{\Rightarrow} (e', \hat{\rho}', \hat{\sigma}', \hat{\kappa}'),$$

which gives us a new iteration function, $\hat{f}'_e : \widehat{System}' \to \widehat{System}'$,

$$\hat{f}'_e(\hat{P}, \hat{\sigma}) = (\hat{P}', \hat{\sigma}'), \text{ where}$$
$$\hat{P}' = \left\{\hat{\pi}' : \hat{\pi} \overset{\hat{\sigma}}{\rightarrowtriangle} (\hat{\pi}', \hat{\sigma}'')\right\} \cup \{\hat{\pi}_0\}$$
$$\hat{\sigma}' = \bigsqcup \left\{\hat{\sigma}'' : \hat{\pi} \overset{\hat{\sigma}}{\rightarrowtriangle} (\hat{\pi}', \hat{\sigma}'')\right\}$$
$$(\hat{\pi}_0, \langle\rangle) = \hat{\mathscr{I}}(e).$$

9.4 Step 4: Pushdown control-flow graphs

Following the earlier graph reformulation of the compact rooted pushdown system, we can reformulate the set of partial configurations as a *pushdown control-flow graph*. A **pushdown control-flow graph** is a frame-action-labeled graph over partial control states, and a **partial control state** is an expression paired with an environment:

$$\widehat{System}'' = \widehat{PDCFG} \times \widehat{Store}$$
$$\widehat{PDCFG} = \mathscr{P}(\widehat{PState}) \times \mathscr{P}(\widehat{PState} \times \widehat{Frame}_{\pm} \times \widehat{PState})$$
$$\hat{\psi} \in \widehat{PState} = \mathsf{Exp} \times \widehat{Env}.$$

In a pushdown control-flow graph, the partial control states are partial configurations which have dropped the continuation component; the continuations are encoded as paths through the graph.

A preliminary analysis of complexity Even without defining the system-space iteration function, we can ask, *How many iterations will it take to reach a fixed point in the worst case?* This question is really asking, *How many edges can we add?* And, *How many entries are there in the store?* Summing these together, we arrive at the worst-case number of iterations:

$$\overbrace{|\widehat{PState}| \times |\widehat{Frame}_{\pm}| \times |\widehat{PState}|}^{\text{PDCFG edges}} + \overbrace{|\widehat{Addr}| \times |\widehat{Clo}|}^{\text{store entries}}.$$

With a monovariant allocation scheme that eliminates abstract environments, the number of iterations ultimately reduces to:

$$|\mathsf{Exp}| \times (2|\widehat{\mathsf{Var}}| + 1) \times |\mathsf{Exp}| + |\mathsf{Var}| \times |\mathsf{Lam}|,$$

which means that, in the worst case, the algorithm makes a cubic number of iterations with respect to the size of the input program.[3]

The worst-case cost of the each iteration would be dominated by a CFL-reachability calculation, which, in the worst case, must consider every state and every edge:

$$O(|\mathsf{Var}|^3 \times |\mathsf{Exp}|^3).$$

Thus, each iteration takes $O(n^6)$ and there are a maximum of $O(n^3)$ iterations, where n is the size of the program. So, total complexity would be $O(n^9)$ for a monovariant pushdown control-flow analysis with this scheme, where n is again the size of the program. Although this algorithm is polynomial-time, we can do better.

9.5 Step 5: Reintroduce ε-closure graphs

Replicating the evolution from Section 8 for this store-widened analysis, we arrive at a more efficient polynomial-time analysis. An ε-closure graph in this setting is a set of pairs of store-less, continuation-less partial states:

$$\widehat{ECG} = \mathscr{P}\left(\widehat{PState} \times \widehat{PState}\right).$$

Then, we can set the system space to include ε-closure graphs:

$$\widehat{System}''' = \widehat{CRPDS} \times \widehat{ECG} \times \widehat{Store}.$$

Before we redefine the iteration function, we need another factored transition relation. The stack- and action-factored transition relation $(\xrightarrow[g]{\hat{\sigma}}) \subseteq \widehat{PState} \times \widehat{PState} \times Store$ determines if a transition is possible under the specified store and stack-action:

$$(e,\hat{\rho}) \xrightarrow[\hat{\phi}_+]{\hat{\sigma}} ((e',\hat{\rho}'),\hat{\sigma}') \text{ iff } (e,\hat{\rho},\hat{\sigma},\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (e',\hat{\rho}',\hat{\sigma}',\hat{\phi}:\hat{\kappa})$$

$$(e,\hat{\rho}) \xrightarrow[\hat{\phi}_-]{\hat{\sigma}} ((e',\hat{\rho}'),\hat{\sigma}') \text{ iff } (e,\hat{\rho},\hat{\sigma},\hat{\phi}:\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (e',\hat{\rho}',\hat{\sigma}',\hat{\kappa})$$

$$(e,\hat{\rho}) \xrightarrow[\varepsilon]{\hat{\sigma}} ((e',\hat{\rho}'),\hat{\sigma}') \text{ iff } (e,\hat{\rho},\hat{\sigma},\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (e',\hat{\rho}',\hat{\sigma}',\hat{\kappa}).$$

Now, we can redefine the iteration function (Figure 7).

Theorem 9.1
Pushdown 0CFA with single-threaded store (PDCFA) can be computed in $O(n^6)$-time, where n is the size of the program.

Proof
As before, the maximum number of iterations is cubic in the size of the program for a monovariant analysis. Fortunately, the cost of each iteration is also now bounded by the number of edges in the graph, which is also cubic. □

[3] In computing the number of frames, we note that in every continuation, the variable and the expression uniquely determine each other based on the let-expression from which they both came. As a result, the number of abstract frames available in a monovariant analysis is bounded by both the number of variables and the number of expressions, *i.e.*, $|\widehat{Frame}| = |\mathsf{Var}|$.

$$\hat{f}_e((\hat{P},\hat{E}),\hat{H},\hat{\sigma}) = ((\hat{P}',\hat{E}'),\hat{H}',\hat{\sigma}''), \text{ where}$$

$$\begin{aligned}
\hat{T}_+ &= \left\{ (\hat{\psi} \overset{\hat{\phi}_+}{\rightarrowtail} \hat{\psi}', \hat{\sigma}') : \hat{\psi} \underset{\hat{\phi}_+}{\overset{\hat{\sigma}}{\rightarrow}} (\hat{\psi}', \hat{\sigma}') \right\} \\
\hat{T}_\varepsilon &= \left\{ (\hat{\psi} \overset{\varepsilon}{\rightarrowtail} \hat{\psi}', \hat{\sigma}') : \hat{\psi} \underset{\varepsilon}{\overset{\hat{\sigma}}{\rightarrow}} (\hat{\psi}', \hat{\sigma}') \right\} \\
\hat{T}_- &= \{ (\hat{\psi}'' \overset{\hat{\phi}_-}{\rightarrowtail} \hat{\psi}''', \hat{\sigma}') : \hat{\psi}'' \underset{\hat{\phi}_-}{\overset{\hat{\sigma}}{\rightarrow}} (\hat{\psi}''', \hat{\sigma}') \text{ and} \\
&\qquad\qquad \hat{\psi} \overset{\hat{\phi}_+}{\rightarrowtail} \hat{\psi}' \in \hat{E} \text{ and} \\
&\qquad\qquad \hat{\psi}' \rightarrowtail \hat{\psi}'' \in \hat{H} \} \\
\hat{T}' &= \hat{T}_+ \cup \hat{T}_\varepsilon \cup \hat{T}_- \\
\hat{E}' &= \{ \hat{e} : (\hat{e}, _) \in \hat{T}' \} \\
\hat{\sigma}'' &= \bigsqcup \{ \hat{\sigma}' : (_, \hat{\sigma}') \in \hat{T}' \} \\
\hat{H}_\varepsilon &= \{ \hat{\psi} \rightarrowtail \hat{\psi}'' : \hat{\psi} \rightarrowtail \hat{\psi}' \in \hat{H} \text{ and } \hat{\psi}' \rightarrowtail \hat{\psi}'' \in \hat{H} \} \\
\hat{H}_{+-} &= \{ \hat{\psi} \rightarrowtail \hat{\psi}''' : \hat{\psi} \overset{\hat{\phi}_+}{\rightarrowtail} \hat{\psi}' \in \hat{E} \text{ and } \hat{\psi}' \rightarrowtail \hat{\psi}'' \in \hat{H} \\
&\qquad\qquad \text{and } \hat{\psi}'' \overset{\hat{\phi}_-}{\rightarrowtail} \hat{\psi}''' \in \hat{E} \} \\
\hat{H}' &= \hat{H}_\varepsilon \cup \hat{H}_{+-} \\
\hat{P}' &= \hat{P} \cup \left\{ \hat{\psi}' : \hat{\psi} \overset{g}{\rightarrowtail} \hat{\psi}' \right\} \cup \{(e, \bot)\}.
\end{aligned}$$

Fig. 7: An ε-closure graph-powered iteration function for pushdown control-flow analysis with a single-threaded store.

10 Introspection for Abstract Garbage Collection

Abstract garbage collection (Might and Shivers 2006b) yields large improvements in precision by using the abstract interpretation of garbage collection to make more efficient use of the finite address space available during analysis. Because of the way abstract garbage collection operates, it grants exact precision to the flow analysis of variables whose bindings die between invocations of the same abstract context. Because pushdown analysis grants exact precision in tracking return-flow, it is clearly advantageous to combine these techniques. Unfortunately, as we shall demonstrate, abstract garbage collection breaks the pushdown model by requiring a full traversal of the stack to discover the root set.

Abstract garbage collection modifies the transition relation to conduct a "stop-and-copy" garbage collection before each transition. To do this, we define a garbage collection function $\hat{G} : \widehat{Conf} \rightarrow \widehat{Conf}$ on configurations:

$$\hat{G}(\overbrace{e, \hat{\rho}, \hat{\sigma}, \hat{\kappa}}^{\hat{c}}) = (e, \hat{\rho}, \hat{\sigma}|_{Reachable(\hat{c}), \hat{\kappa}}),$$

where the pipe operation $f|S$ yields the function f, but with inputs not in the set S mapped to bottom—the empty set. The reachability function $\mathit{Reachable} : \widehat{\mathit{Conf}} \to \mathscr{P}(\widehat{\mathit{Addr}})$ first computes the root set, and then the transitive closure of an address-to-address adjacency relation:

$$\mathit{Reachable}(\overbrace{e, \hat{\rho}, \hat{\sigma}, \hat{\kappa}}^{\hat{c}}) = \left\{ \hat{a} : \hat{a}_0 \in \mathit{Root}(\hat{c}) \text{ and } \hat{a}_0 \underset{\hat{\sigma}}{\overset{*}{\rightarrowtriangle}} \hat{a} \right\},$$

where the function $\mathit{Root} : \widehat{\mathit{Conf}} \to \mathscr{P}(\widehat{\mathit{Addr}})$ finds the root addresses:

$$\mathit{Root}(e, \hat{\rho}, \hat{\sigma}, \hat{\kappa}) = \mathit{range}(\hat{\rho}) \cup \mathit{StackRoot}(\hat{\kappa}),$$

and the $\mathit{StackRoot} : \widehat{\mathit{Kont}} \to \mathscr{P}(\widehat{\mathit{Addr}})$ function finds roots down the stack:

$$\mathit{StackRoot}\langle \phi_1, \ldots, \phi_n \rangle = \bigcup_i \mathscr{T}(\phi_i),$$

using a "touches" function, $\mathscr{T} : \widehat{\mathit{Frame}} \to \mathscr{P}(\widehat{\mathit{Addr}})$:

$$\mathscr{T}(v, e, \hat{\rho}) = \mathit{range}(\hat{\rho}),$$

and the relation $(\rightarrowtriangle) \subseteq \widehat{\mathit{Addr}} \times \widehat{\mathit{Store}} \times \widehat{\mathit{Addr}}$ connects adjacent addresses:

$$\hat{a} \underset{\hat{\sigma}}{\rightarrowtriangle} \hat{a}' \text{ iff there exists } (\mathit{lam}, \hat{\rho}) \in \hat{\sigma}(\hat{a}) \text{ such that } \hat{a}' \in \mathit{range}(\hat{\rho}).$$

The new abstract transition relation is thus the composition of abstract garbage collection with the old transition relation:

$$(\widehat{\Rightarrow}_{\mathrm{GC}}) = (\widehat{\Rightarrow}) \circ \hat{G}.$$

Problem: Stack traversal violates pushdown constraint In the formulation of pushdown systems, the transition relation is restricted to looking at the top frame, and in less restricted formulations that may read the stack, the reachability decision procedures need the entire system up-front. Thus, the relation $(\widehat{\Rightarrow}_{\mathrm{GC}})$ cannot be computed as a straightforward pushdown analysis using summarization.

Solution: Introspective pushdown systems To accommodate the richer structure of the relation $(\widehat{\Rightarrow}_{\mathrm{GC}})$, we now define *introspective* pushdown systems. Once defined, we can embed the garbage-collecting abstract interpretation within this framework, and then focus on developing a control-state reachability algorithm for these systems.

An **introspective pushdown system** is a quadruple $M = (Q, \Gamma, \delta, q_0)$:

1. Q is a finite set of control states;
2. Γ is a stack alphabet;
3. $\delta \subseteq Q \times \Gamma^* \times \Gamma_\pm \times Q$ is a transition relation (where $(q, \kappa, \phi_-, q') \in \delta$ implies $\kappa \equiv \phi : \kappa'$); and
4. q_0 is a distinguished root control state.

The second component in the transition relation is a realizable stack at the given control-state. This realizable stack distinguishes an introspective pushdown system from a general pushdown system. $\mathbb{IPDS}$ denotes the class of all introspective pushdown systems.

Determining how (or if) a control state q transitions to a control state q', requires knowing a path taken to the state q. We concern ourselves with root-reachable states. When $M = (Q, \Gamma, \delta, q_0)$, if there is a $\hat{\kappa}$ such that $(q_0, \langle\rangle) \underset{M}{\longmapsto}^* (q, \hat{\kappa})$ we say q is reachable via $\hat{\kappa}$, where

$$\frac{}{(q, \hat{\kappa}) \underset{M}{\longmapsto}^* (q, \hat{\kappa})} \qquad \frac{(q, \hat{\kappa}) \underset{M}{\longmapsto}^* (q', \hat{\kappa}') \quad (q', \hat{\kappa}', g', q'') \in \delta}{(q, \hat{\kappa}) \underset{M}{\longmapsto}^* (q'', [\hat{\kappa}'_+ g'])}$$

10.1 Garbage collection in introspective pushdown systems

To convert the garbage-collecting, abstracted CESK machine into an introspective pushdown system, we use the function $\widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}} : \mathsf{Exp} \to \mathbb{IPDS}$:

$$\begin{aligned} \widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}}(e) &= (Q, \Gamma, \delta, q_0) \\ Q &= \mathsf{Exp} \times \widehat{Env} \times \widehat{Store} \\ \Gamma &= \widehat{Frame} \\ (q_0, \langle\rangle) &= \hat{\mathscr{I}}(e) \end{aligned}$$

$$\begin{aligned} (q, \hat{\kappa}, \varepsilon, q') \in \delta \text{ iff } \hat{G}(q, \hat{\kappa}) &\mathrel{\widehat{\Rightarrow}} (q', \hat{\kappa}) \\ (q, \hat{\phi} : \hat{\kappa}, \hat{\phi}_-, q') \in \delta \text{ iff } \hat{G}(q, \hat{\phi} : \hat{\kappa}) &\mathrel{\widehat{\Rightarrow}} (q', \hat{\kappa}) \\ (q, \hat{\kappa}, \hat{\phi}_+, q') \in \delta \text{ iff } \hat{G}(q, \hat{\kappa}) &\mathrel{\widehat{\Rightarrow}} (q', \hat{\phi} : \hat{\kappa}). \end{aligned}$$

11 Problem: Reachability for Introspective Pushdown Systems is Uncomputable

As currently formulated, computing control-state reachability for introspective pushdown systems is uncomputable. The problem is that the transition relation expects to enumerate every possible stack for every control point at every transition, without restriction.

Theorem 11.1
Reachability in introspective pushdown systems is uncomputable.

Proof
Consider an IPDS with two states — `searching` (start state) and `valid` — and a singleton stack alphabet of unit ($\top$). For any first-order logic proposition, ϕ, we can define a reduction relation that interprets the length of the stack as an encoding of a proof of ϕ. If the length encodes an ill-formed proof object, or is not a proof of ϕ, `searching` pushes $\top$ on the stack and transitions to itself. If the length encodes a proof of ϕ, transition to `valid`. By the completeness of first-order logic, if ϕ is valid, there is a finite proof, making the pushdown system terminate in `valid`. If it is not valid, then there is no proof and `valid` is unreachable. Due to the undecidability of first-order logic, we definitely cannot have a decision procedure for reachability of IPDSs. □

To make introspective pushdown systems computable, we must first refine our definition of introspective pushdown systems to operate on *sets* of stacks and insist these sets be regular.

A **conditional pushdown system** (CPDS) is a quadruple $M = (Q, \Gamma, \delta, q_0)$:

1. Q is a finite set of control states;
2. Γ is a stack alphabet;
3. $\delta \subseteq_{\text{fin}} Q \times \mathscr{REG}(\Gamma^*) \times \Gamma_\pm \times Q$ is a transition relation (same restriction on stacks); and
4. q_0 is a distinguished root control state,

where $\mathscr{REG}(S)$ is the set of all regular languages formable with strings in S.

The regularity constraint on the transition relations guarantees that we can decide applicability of transition rules at each state, since (as we will see) the set of all stacks that reach a state in a CPDS has decidable overlap with regular languages. Let $\mathbb{CPDS}$ denote the set of all conditional pushdown systems.

The rules for reachability with respect to sets of stacks are similar to those for IPDSs.

$$\frac{}{(q,\hat{\kappa}) \underset{M}{\longmapsto}^* (q,\hat{\kappa})} \qquad \frac{(q,\hat{\kappa}) \underset{M}{\longmapsto}^* (q',\hat{\kappa}') \quad \hat{\kappa}' \in \hat{K}' \quad (q',\hat{K}',g',q'') \in \delta}{(q,\hat{\kappa}) \underset{M}{\longmapsto}^* (q'', [\hat{\kappa}'_+ g'])}$$

We will write $q \overset{\hat{K},g}{\underset{M}{\longmapsto\!\!\!\!\rightarrow}} q'$ to mean there are $\hat{\kappa}, \hat{K}$ such that q is reachable via $\hat{\kappa}$, $\hat{\kappa} \in \hat{K}$ and $(q,\hat{K},g,q') \in \delta$. We will omit the labels above if they merely exist.

11.1 Garbage collection in conditional pushdown systems

Of course, we must adapt abstract garbage collection to this refined framework. To convert the garbage-collecting, abstracted CESK machine into a conditional pushdown system, we use the function $\widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}'} : \mathsf{Exp} \to \mathbb{CPDS}$:

$$\widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}'}(e) = (Q, \Gamma, \delta, q_0)$$
$$Q = \mathsf{Exp} \times \widehat{Env} \times \widehat{Store}$$
$$\Gamma = \widehat{Frame}$$

For all sets of addresses $A \subseteq \widehat{Addr}$ let $\hat{K} = \{\hat{\kappa} : \mathit{StackRoot}(\hat{\kappa}) = A\}$

$$(q,\hat{K},\varepsilon,q') \in \delta \text{ iff } \hat{G}(q,\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (q',\hat{\kappa}) \text{ for any } \hat{\kappa} \in \hat{K}$$
$$(q,\hat{K},\hat{\phi}_-,q') \in \delta \text{ iff } \hat{G}(q,\hat{\phi} : \hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (q',\hat{\kappa}) \text{ for any } \hat{\phi} : \hat{\kappa} \in \hat{K}$$
$$(q,\hat{K},\hat{\phi}_+,q') \in \delta \text{ iff } \hat{G}(q,\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (q',\hat{\phi} : \hat{\kappa}) \text{ for any } \hat{\kappa} \in \hat{K}$$
$$(q_0, \langle\rangle) = \hat{\mathscr{I}}(e).$$

Assuming we can overcome the difficulty of computing with some representation of a set of stacks, the intuition for the decidability of control-state reachability with garbage collection stems from two observations: garbage collection operates on sets of addresses, and for any given control point there is a finite number of sets of sets of addresses. The finiteness makes the definition of δ fit the finiteness restriction of CPDSs. The regularity of $\hat{K}$ (for any given A, which we recall are finite sets) is apparent from a simple construction: let the DFA control states represent the subsets of A, with $\emptyset$ the start state and A the accepting state. Transition from $A' \subseteq A$ to $A \cup \mathscr{T}(\hat{\phi})$ for each $\hat{\phi}$ (no transition if the result

is not a subset of A). Thus any string of frames that has a stack root of A (and only A) gets accepted.

The last challenge to consider before we can delve into the mechanics of computing reachable control states is *how* to represent the sets of stacks that may be paired with each control state. Fortunately, a regular language can describe the stacks that share the same root addresses, the set of stacks at a control point are recognized by a one-way non-deterministic stack automaton (1NSA), *and*, fortuitously, non-empty overlap of these two is decidable (but NP-hard (Rounds 1973)). The 1NSA describing the set of stacks at a control point is already encoded in the structure of the (augmented) CRPDS that we will accumulate while computing reachable control states. As we develop an algorithm for control-state reachability, we will exploit this insight (Section 13).

12 Reachability in Conditional Pushdown Systems

We will show a progression of constructions that take us along the following line:

$$\mathbb{CPDS} \xrightarrow[\S 12.2]{} \mathbb{CCPDS} \xrightarrow[\S 12.3]{\text{specialize}} \textit{PDCFA with GC} \rightarrow \textit{approx. PDCFA with GC}$$

In the first construction, we show that a CCPDS is finitely constructible in a similar fashion as in Section 7. The key is to take the current introspective CRPDS and "read off" an automaton that describes the stacks accepted at each state. For traditional pushdown systems, this is always an NFA, but introspection adds another feature: transition if the string accepted so far is accepted by a given NFA. Such power falls outside of standard NFAs and into one-way non-deterministic stack automata (1NSA)[4]. These automata enjoy closure under finite intersection with regular languages and decidable emptiness checking (Ginsburg et al. 1967), which we use to decide applicability of transition rules. If the stacks realizable at q have a non-empty intersection with a set of stacks $\hat{K}$ in a rule $(q, \hat{K}, g, q') \in \delta$, then there are paths from the start state to q that further reach q'.

The structure of the GC problem allows us to sidestep the 1NSA constructions and more directly compute state reachability. We specialize to garbage collection in subsection 12.3. We finally show a space-saving approximation that our implementation uses.

12.1 One-way non-deterministic stack automata

The machinery we use for describing the realizable stacks at a state is a generalized pushdown automaton itself. A stack automaton is permitted to move a cursor up and down the stack and read frames (left and right on the input if two-way, only right if one-way), but only push and pop when the stack cursor is at the top. Formally, a **one-way stack automaton** is a 6-tuple $A = (Q, \Sigma, \Gamma, \delta, q_0, F)$ where

1. Q is a finite nonempty set of states,
2. Σ is a finite nonempty input alphabet,
3. Γ is a finite nonempty stack alphabet,
4. $\delta \subseteq Q \times (\Gamma \cup \{\varepsilon\}) \times (\Sigma \cup \{\varepsilon\}) \times \{\uparrow, \cdot, \downarrow\} \times \Gamma_{\pm} \times Q$ is the transition relation,

[4] The reachable states of a 1NSA is known to be regular, but the paths are not.

5. $q_0 \in Q$ is the start state, and
6. $F \subseteq Q$ the set of final states

An element of the transition relation, $(q, \phi_\varepsilon, a, d, \phi_\pm, q')$, should be read as, "if at q the right of the stack cursor is prefixed by ϕ_ε and the input is prefixed by a, then consume a of the input, transition to state q', move the stack cursor in direction d, and if at the top of the stack, perform stack action $\phi_\pm$." This reading translates into a run relation on **instantaneous descriptions**, $Q \times (\Gamma^* \times \Gamma^*) \times \Sigma^*$. These descriptions are essentially machine states that hold the current control state, the stack split around the cursor, and the rest of the input.

$$\frac{(q, \phi_\varepsilon, a, d, \phi_\pm, q') \in \delta \quad \phi_\varepsilon \sqsubseteq \Gamma_T \quad w \equiv aw' \quad [\Gamma'_B, \Gamma'_T] = P(\phi_\pm, D(d, [\Gamma_B, \Gamma_T]))}{(q, [\Gamma_B, \Gamma_T], w) \longmapsto (q', [\Gamma'_B, \Gamma'_T], w')}$$

where

$$\begin{aligned} P(\phi_+, [\Gamma_B, \phi'_\varepsilon]) &= [\Gamma_B \phi'_\varepsilon, \phi] \\ P(\phi_-, [\Gamma_B \phi', \phi]) &= [\Gamma_B, \phi'] \\ P(\phi_-, [\varepsilon, \phi]) &= [\varepsilon, \varepsilon] \\ P(\phi_\pm, \Gamma_{B,T}) &= \Gamma_{B,T} \quad \text{otherwise} \end{aligned}$$

$$\begin{aligned} D(\uparrow, [\Gamma_B, \phi\Gamma_T]) &= [\Gamma_B \phi, \Gamma_T] \\ D(\downarrow, [\Gamma_B \phi, \Gamma_T]) &= [\Gamma_B, \phi\Gamma_T] \\ D(d, \Gamma_{B,T}) &= \Gamma_{B,T} \quad \text{otherwise} \end{aligned}$$

The meta-functions P and D perform the stack actions and direct the stack cursor, respectively. A string w is thus accepted by a 1NSA A iff there are $q \in F, \Gamma_B, \Gamma_T \in \Gamma^*$ such that

$$(q_0, [\varepsilon, \varepsilon], w) \longmapsto^* (q, [\Gamma_B, \Gamma_T], \varepsilon)$$

Next we develop an introspective form of compact rooted pushdown systems that use 1NSAs for realizable stacks, and prove a correspondence with conditional pushdown systems.

12.2 Compact conditional pushdown systems

Similar to rooted pushdown systems, we say a conditional pushdown system $G = (S, \Gamma, E, s_0)$ is compact if all states, frames and edges are on some path from the root. We will refer to this class of conditional pushdown systems as $\mathbb{CCPDS}$. Assuming we have a way to decide overlap between the set of realizable stacks at a state and a regular language of stacks, we can compute the CCPDS in much the same way as in section 7.

$$\begin{aligned} &\mathscr{F}(M) = f, \text{ where} \\ &\quad M = (Q, \Gamma, \delta, q_0) \\ &\overbrace{f(S, \Gamma, E, s_0)}^{G} = (S', \Gamma, E', s_0), \text{ where} \\ &\quad S' = S \cup \left\{ s' : s \in S \text{ and } s \underset{M}{\longmapsto\!\!\!\!\!\rightarrow} s' \right\} \cup \{s_0\} \\ &\quad E' = E \cup \left\{ s \overset{\hat{K}, g}{\longmapsto} s' : s \in S \text{ and } s \underset{M}{\overset{\hat{K}, g}{\longmapsto\!\!\!\!\!\rightarrow}} s' \text{ and } \mathit{Stacks}(G)(s) \cap \hat{K} \neq \emptyset \right\}. \end{aligned}$$

The function $\mathit{Stacks} : \mathbb{CCPDS} \to S \to 1\mathbb{NSA}$ performs the stack extraction with a construction that inserts the stack-checking NFA for each reduction rule after it has run the cursor to the bottom of the stack, and continues from the final states to the state dictated by the rule (added by meta-function *gadget*). All the stack manipulations from s_0 to s are ε-transitions in terms of reading input; only once control reaches s do we check if the stack is the same as the input, which captures the notion of a stack realizable at s. Once control reaches s, we run down to the bottom of the stack again, and then match the stack against the input; complete matches are accepted. To determine the bottom and top of the stack, we add distinct sentinel symbols to the stack alphabet, $\mathfrak{c}$ and \$.

$$\mathit{Stacks}(\overbrace{S,\Gamma,E,s_0}^{G})(s) = (S \cup S', \Gamma, \Gamma \cup \{\mathfrak{c},\$\}, \delta, s_{\text{start}}, \{s_{\text{final}}\}), \text{ where}$$

$s_{\text{start}}, s_{\text{down}}, s_{\text{check}}, s_{\text{final}}$ fresh, and S', δ the smallest sets such that

$$\{s_{\text{start}}, s_{\text{down}}, s_{\text{check}}, s_{\text{final}}\} \subseteq S'$$
$$(s_{\text{start}}, \varepsilon, \varepsilon, \cdot, \mathfrak{c}_+, s_0) \in \delta$$
$$\mathit{gadget}(s', \hat{K}, \gamma_\pm, s'') \sqsubseteq (\delta, S') \text{ if } (s', \hat{K}, \gamma_\pm, s'') \in E$$
$$(s, \varepsilon, \varepsilon, \cdot, \$_+, s_{\text{down}}) \in \delta$$
$$(s_{\text{down}}, \varepsilon, \varepsilon, \downarrow, \varepsilon, s_{\text{down}}) \in \delta$$
$$(s_{\text{down}}, \mathfrak{c}, \varepsilon, \uparrow, \varepsilon, s_{\text{check}}) \in \delta$$
$$(s_{\text{check}}, a, a, \uparrow, \varepsilon, s_{\text{check}}) \in \delta,\ a \in \Gamma \cup \{\varepsilon\}$$
$$(s_{\text{check}}, \$, \varepsilon, \uparrow, \varepsilon, s_{\text{final}}) \in \delta$$

The first rule changes the initial state to initialize the stack with the "bottom" sentinel. Every reduction of the CPDS is given the gadget discussed above and explained below. The last five rules are what implement the final checking of stack against input. When at the state we are recognizing realizable stacks for, the machine will have the cursor at the top of the stack, so we push the "top" sentinel before moving the cursor all the way down to the bottom. When s_{down} finds the bottom sentinel at the cursor, it moves the cursor past it to start the exact matching in s_{check}. If the cursor matches the input exactly, we consume the input and move the cursor past the matched character to start again. When s_{check} finds the top sentinel, it transitions to the final state; if the input is not completely exhausted, the machine will get stuck and not accept.

$$
\begin{aligned}
&\mathit{gadget}(s,\hat{K},\gamma_{\pm},s') = (\delta', Q \cup \{q_{\text{down}}, q_{\text{out}}\}) \text{ where}\\
&\text{Let } N = (Q,\Sigma,\delta,q_0,F) \text{ be a fresh NFA recognizing } \hat{K},\ q_{\text{down}}, q_{\text{out}} \text{ fresh states}\\
&(q,a,\varepsilon,\uparrow,\varepsilon,q') \in \delta' \text{ if } (q,a,q') \in \delta,\ a \in \Sigma\\
&(q,\varepsilon,\varepsilon,\cdot,\varepsilon,q') \in \delta' \text{ if } (q,\varepsilon,q') \in \delta\\
&(q,\$,\varepsilon,\cdot,\$_{-},q_{\text{out}}) \in \delta' \text{ if } q \in F\\
&(q_{\text{out}},\varepsilon,\varepsilon,\cdot,\gamma_{\pm},s') \in \delta'\\
&(s,\varepsilon,\varepsilon,\cdot,\$_{+},q_{\text{down}}) \in \delta'\\
&(q_{\text{down}},\varepsilon,\varepsilon,\downarrow,\varepsilon,q_{\text{down}}) \in \delta'\\
&(q_{\text{down}},\not\varepsilon,\varepsilon,\uparrow,\varepsilon,q_0) \in \delta'
\end{aligned}
$$

We explain each rule in order. When the NFA that recognizes $\hat{K}$ consumes a character, the stack automaton should similarly read the character on the stack and move the cursor along. If the NFA makes an ε-transition, the stack automaton should also, without moving the stack cursor. When this sub-machine N is in a final state, the cursor should be at the top of the stack (if indeed it matched), so we pop off the top sentinel and proceed to do the stack action the IPDS does when transitioning to the next state. The last three rules implement the same "run down to the bottom" gadget used before, when matching the stack against the input.

Finally, we can show that states are reachable in a conditional pushdown system iff they are reached in their corresponding CCPDS. Consider a map

$$\mathscr{CC} : \mathbb{CPDS} \to \mathbb{CCPDS}$$

such that given a conditional pushdown system $M = (Q,\Gamma,\delta,q_0)$, its equivalent CCPDS is $\mathscr{CC}(M) = (S,\Gamma,E,q_0)$ where S contains reachable nodes:

$$S = \left\{ q : (q_0,\langle\rangle) \underset{M}{\longmapsto}^{*} (q,\hat{\kappa}) \right\}$$

and the set E contains reachable edges:

$$E = \left\{ q \overset{\hat{K},g}{\rightarrowtail} q' : q \underset{M}{\overset{\hat{K},g}{\longmapsto\!\!\!\!\!\rightarrow}} q' \right\}$$

Theorem 12.1 (*Computable reachability*)
For all $M \in \mathbb{CPDS}$, $\mathscr{CC}(M) = \text{lfp}(\mathscr{F}(M))$

Proof in Appendix 18.2.

Corollary 12.1 (*Realizable stacks of CPDSs are recognized by 1NSAs*)
For all $M = (Q,\Gamma,\delta,q_0) \in \mathbb{CPDS}$, and $(S,\Gamma,E,q_0) = \text{lfp}(\mathscr{F}(M))$, $(q_0,\langle\rangle) \underset{M}{\longmapsto}^{*} (q,\hat{\kappa})$ iff $q \in S$ and $\mathit{Stacks}(G)(q)$ accepts $\hat{\kappa}$.

12.3 Simplifying garbage collection in conditional pushdown systems

The decision problems on 1NSAs are computationally intractable in general, but luckily GC is a special problem where we do not need the full power of 1NSAs. There are equally

$$\hat{f}_e((\hat{P},\hat{E}),\hat{H}) = ((\hat{P}',\hat{E}'),\hat{H}'), \text{ where}$$

$$\hat{E}_+ = \left\{ (\hat{\psi},A) \overset{\hat{\phi}_+}{\rightarrowtail} (\hat{\psi}',A \cup \mathscr{T}(\hat{\phi})) : \hat{\psi} \underset{\hat{\phi}_+}{\overset{A}{\rightarrow}} \hat{\psi}' \right\}$$

$$\hat{E}_\varepsilon = \left\{ (\hat{\psi},A) \overset{\varepsilon}{\rightarrowtail} (\hat{\psi}',A) : \hat{\psi} \underset{\varepsilon}{\overset{A}{\rightarrow}} \hat{\psi}' \right\}$$

$$\hat{E}_- = \left\{ (\hat{\psi}'',A) \overset{\hat{\phi}_-}{\rightarrowtail} (\hat{\psi}''',A') : \hat{\psi}'' \underset{\hat{\phi}_-}{\overset{A}{\rightarrow}} \hat{\psi}''' \text{ and } (\hat{\psi},A') \overset{\hat{\phi}_+}{\rightarrowtail} (\hat{\psi}',A) \in \hat{E} \text{ and } (\hat{\psi}',A) \rightarrowtail (\hat{\psi}'',A) \in \hat{H} \right\}$$

$$\hat{E}' = \hat{E}_+ \cup \hat{E}_\varepsilon \cup \hat{E}_-$$

$$\hat{H}_\varepsilon = \left\{ \hat{\Omega} \rightarrowtail \hat{\Omega}'' : \hat{\Omega} \rightarrowtail \hat{\Omega}' \in \hat{H} \text{ and } \hat{\Omega}' \rightarrowtail \hat{\Omega}'' \in \hat{H} \right\}$$

$$\hat{H}_{+-} = \left\{ \hat{\Omega} \rightarrowtail \hat{\Omega}''' : \hat{\Omega} \overset{\hat{\phi}_+}{\rightarrowtail} \hat{\Omega}' \in \hat{E} \text{ and } \hat{\Omega}' \rightarrowtail \hat{\Omega}'' \in \hat{H} \right.$$
$$\left. \text{and } \hat{\Omega}'' \overset{\hat{\phi}_-}{\rightarrowtail} \hat{\Omega}''' \in \hat{E} \right\}$$

$$\hat{H}' = \hat{H}_\varepsilon \cup \hat{H}_{+-}$$

$$\hat{P}' = \hat{P} \cup \left\{ \hat{\Omega}' : \hat{\Omega} \overset{g}{\rightarrowtail} \hat{\Omega}' \right\} \cup \{((e,\bot,\bot),\emptyset)\}.$$

Fig. 8: An ε-closure graph-powered iteration function for pushdown garbage-collecting control-flow analysis

precise techniques at much lower cost, and less precise techniques that can shrink the explored state space.[5] The transition relation we build does not enumerate all sets of addresses, but instead queries the graph for the sets of addresses it should consider in order to apply GC. A fully precise method to manage the stack root addresses is to add the root addresses to the representation of each state, and update it incrementally. The root addresses can be seen as the representation of $\hat{K}$ in edge labels, but to maintain the precision, the set must also distinguish control states. This addition to the state space is an effective *reification* of the stack filtering that conditional performs.

An approximative method is to not distinguish control states, but rather to traverse the graph backward through ε-closure edges and push edges, and collect the root addresses through the pushed frames. As more paths are discovered to control states, more stacks will be realizable there, which add more to the stack root addresses to consider as the relation steps forward. For soundness, edges still must be labeled with the language of stacks they are valid for, since they can become invalid as more stacks reach control states. Notice that the root sets of addresses are isomorphic to languages of stacks that have the given root set, so we can use sets of addresses as the language representation.

We consider both methods in turn, augmenting the compaction algorithm from subsection 9.5. Each have program states that consist of the expression, environment, *and* store; $\hat{\psi} \in \widehat{PState} = \mathsf{Exp} \times \widehat{Env} \times \widehat{Store}$. Since GC is a non-monotonic operation, stores cannot

[5] The added precision of GC with tighter working sets makes the state space comparison between the two approaches non-binary. Neither approach is clearly better in terms of performance.

be shared globally without sacrificing the precision benefits of GC. For the first method, program states additionally include the stack root set of addresses; we will call these ornamented program states, $\hat{\Omega} \in \widehat{OPState} = \widehat{PState} \times \mathscr{P}(\widehat{Addr})$. We show the non-worklist solution to computing reachability by employing the function $\hat{f}_e$ defined in Figure 8. In order to define the iteration function, we need a refactored transition relation $\xrightarrow[g]{A} \subseteq \widehat{PState} \times \widehat{PState}$, defined as follows:

$$(e,\hat{\rho},\hat{\sigma}) \xrightarrow[\hat{\phi}_+]{A} (e',\hat{\rho}',\hat{\sigma}') \text{ iff } \mathit{StackRoot}(\hat{\kappa}) = A \text{ and } \hat{G}(e,\hat{\rho},\hat{\sigma},\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (e',\hat{\rho}',\hat{\sigma}',\hat{\phi}:\hat{\kappa})$$

$$(e,\hat{\rho},\hat{\sigma}) \xrightarrow[\hat{\phi}_-]{A} (e',\hat{\rho}',\hat{\sigma}') \text{ iff } \mathit{StackRoot}(\hat{\phi}:\hat{\kappa}) = A \text{ and } \hat{G}(e,\hat{\rho},\hat{\sigma},\hat{\phi}:\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (e',\hat{\rho}',\hat{\sigma}',\hat{\kappa})$$

$$(e,\hat{\rho},\hat{\sigma}) \xrightarrow[\varepsilon]{A} (e',\hat{\rho}',\hat{\sigma}') \text{ iff } \mathit{StackRoot}(\hat{\kappa}) = A \text{ and } \hat{G}(e,\hat{\rho},\hat{\sigma},\hat{\kappa}) \mathrel{\widehat{\Rightarrow}} (e',\hat{\rho}',\hat{\sigma}',\hat{\kappa})$$

Theorem 12.2 (*Correctness of GC specialization*)
$\mathrm{lfp}(\hat{f}_e)$ completely abstracts $\widehat{\mathscr{C}\mathscr{C}}(\widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}'}(e))$

The approximative method changes the representation of edges in the graph to contain sets of addresses, $\hat{E} \in \widehat{Edge} = \mathscr{P}\left(\widehat{PState} \times \mathscr{P}(\widehat{Addr}) \times \widehat{Frame}_{\pm} \times \widehat{PState}\right)$. We also add in a sub-fixed-point computation for $\hat{t} : (\widehat{PState} \to \mathscr{P}(\widehat{Addr})) \to (\widehat{PState} \to \mathscr{P}(\widehat{Addr}))$, to traverse the graph and collect the union of all stack roots for stacks realizable at a state. Although we show a non-worklist solution here (in Figure 9) to not be distracting, this solution will not compute the same reachable states as a worklist solution due to the ever-growing stack roots at each state. Only states in the worklist would need to be analyzed at the larger stack root sets. In other words, the non-worklist solution potentially throws in more live addresses at states that would otherwise not need to be re-examined.

This approximation is not an easily described introspective pushdown system since the root sets it uses depend on the iteration state — particularly what frames have reached a state *so far*, regardless of the stack filtering the original CPDS performs. The regular sets of stacks acceptable at some state can be extracted *a posteriori* from the fixed point of the function $\hat{f}'_e$ defined in Figure 9, if so desired. The next theorem follows from the fact that $\mathscr{R}(\hat{\psi}) \supseteq A$ for any represented $(\hat{\psi},A)$.

Theorem 12.3 (*Approximate GC is sound*)
$\mathrm{lfp}(\hat{f}'_e)$ approximates $\mathrm{lfp}(\hat{f}_e)$.

The last thing to notice is that by disregarding the filtering, the stack root set can get larger and render previous GCs unsound, since more addresses can end up live than were previously considered. Thus we label edges with the root set for which the GC was considered, in order to not make false predictions.

13 Implementing Introspective Pushdown Analysis with Garbage Collection

The reachability-based analysis for a pushdown system described in the previous section requires two mutually-dependent pieces of information in order to add another edge:

$$\hat{f}'_e((\hat{P},\hat{E}),\hat{H}) = ((\hat{P}',\hat{E}'),\hat{H}'), \text{ where}$$

$$\begin{aligned}
\hat{E}_+ &= \left\{ \hat{\psi} \underset{\hat{\phi}_+}{\overset{A}{\rightarrowtail}} \hat{\psi}' : A = \mathscr{R}(\hat{\psi}), \hat{\psi} \underset{\hat{\phi}_+}{\overset{A}{\rightarrow}} \hat{\psi}' \right\} \\
\hat{E}_\varepsilon &= \left\{ \hat{\psi} \underset{\varepsilon}{\overset{A}{\rightarrowtail}} \hat{\psi}' : A = \mathscr{R}(\hat{\psi}), \hat{\psi} \underset{\varepsilon}{\overset{A}{\rightarrow}} \hat{\psi}' \right\} \\
\hat{E}_- &= \{ \hat{\psi}'' \underset{\hat{\phi}_-}{\overset{A}{\rightarrowtail}} \hat{\psi}''' : A = \mathscr{R}(\hat{\psi}''), \hat{\psi}'' \underset{\hat{\phi}_-}{\overset{A}{\rightarrow}} \hat{\psi}''' \text{ and} \\
&\qquad \hat{\psi} \underset{\hat{\phi}_+}{\overset{A'}{\rightarrowtail}} \hat{\psi}' \in \hat{E} \text{ and} \\
&\qquad \hat{\psi}' \rightarrowtail \hat{\psi}'' \in \hat{H} \} \\
\hat{E}' &= \hat{E}_+ \cup \hat{E}_\varepsilon \cup \hat{E}_- \\
\hat{H}_\varepsilon &= \{ \hat{\psi} \rightarrowtail \hat{\psi}'' : \hat{\psi} \rightarrowtail \hat{\psi}' \in \hat{H} \text{ and } \hat{\psi}' \rightarrowtail \hat{\psi}'' \in \hat{H} \} \\
\hat{H}_{+-} &= \{ \hat{\psi} \rightarrowtail \hat{\psi}''' : \hat{\psi} \underset{\hat{\phi}_+}{\overset{A}{\rightarrowtail}} \hat{\psi}' \in \hat{E} \text{ and } \hat{\psi}' \rightarrowtail \hat{\psi}'' \in \hat{H} \\
&\qquad \text{and } \hat{\psi}'' \underset{\hat{\phi}_-}{\overset{A'}{\rightarrowtail}} \hat{\psi}''' \in \hat{E} \} \\
\hat{H}' &= \hat{H}_\varepsilon \cup \hat{H}_{+-} \\
\hat{P}' &= \hat{P} \cup \left\{ \hat{\psi}' : \hat{\psi} \underset{g}{\overset{A}{\rightarrowtail}} \hat{\psi}' \right\} \cup \{(e,\bot,\bot)\} \\
\hat{t}(\mathscr{R}) &= \lambda\, \hat{\psi}. \bigcup (\left\{ \mathscr{T}(\hat{\phi}) \cup \mathscr{R}(\hat{\psi}') : \hat{\psi}' \underset{\hat{\phi}_+}{\overset{A}{\rightarrowtail}} \hat{\psi} \in \hat{E} \right\} \cup \{ \mathscr{R}(\hat{\psi}') : \hat{\psi}' \rightarrowtail \hat{\psi} \in \hat{H} \}) \\
\mathscr{R} &= \mathrm{lfp}(\hat{t}).
\end{aligned}$$

Fig. 9: Approximate pushdown garbage-collecting control-flow analysis.

1. The topmost frame on a stack for a given control state q. This is essential for *return* transitions, as this frame should be popped from the stack and the store and the environment of a caller should be updated respectively.
2. Whether a given control state q is reachable or not from the initial state q_0 along realizable sequences of stack actions. For example, a path from q_0 to q along edges labeled "push, pop, pop, push" is not realizable: the stack is empty after the first pop, so the second pop cannot happen—let alone the subsequent push.

Knowing about a possible topmost frame on a stack and initial-state reachability is enough for a classic pushdown reachability summarization to proceed one step further, and we presented an efficient algorithm to compute those in Section 8. However, to deal with the presence of an abstract GC in a *conditional* PDS, we add:

3. For a given control state q, what are the touched addresses of *all* possible frames that could happen to be *on* the stack at the moment the CPDS is in the state q?

The crucial addition to the algorithm is maintaining for each node q' in the CRPDS a set of ε-*predecessors*, i.e., nodes q, such that $q \xmapsto[M]{\vec{g}}\!\!\!\!\rightarrow\!\!\!\!\rightarrow q'$ and $[\vec{g}] = \varepsilon$. In fact, only two out of three kinds of transitions can cause a change to the set of ε-predecessors for a particular node q: an addition of an ε-edge or a pop edge to the CRPDS.

One can notice a subtle mutual dependency between computation of ε-predecessors and top frames during the construction of a CRPDS. Informally:

- A *top frame* for a state q can be pushed as a direct predecessor (*e.g.*, q follows a nested let-binding), or as a direct predecessor to an ε-predecessor (*e.g.*, q is in tail position and will return to a waiting let-binding).
- When a new ε-edge $q \xrightarrow{\varepsilon} q'$ is added, all ε-predecessors of q become also ε-predecessors of q'. That is, ε-summary edges are transitive.
- When a γ_--pop-edge $q \xrightarrow{\gamma_-} q'$ is added, new ε-predecessors of a state q_1 can be obtained by checking if q' is an ε-predecessor of q_1 and examining all existing ε-predecessors of q, such that γ_+ is their possible top frame: this situation is similar to the one depicted in the example above.

The third component—the touched addresses of *all* possible frames on the stack for a state q—is straightforward to compute with ε-predecessors: starting from q, trace out only the edges which are labeled ε (summary or otherwise) or γ_+. The frame for any action γ_+ in this trace is a possible stack action. Since these sets grow monotonically, it is easy to cache the results of the trace, and in fact, propagate incremental changes to these caches when new ε-summary or γ_+ nodes are introduced. This implementation strategy captures the approximative approach to performing GC, as discussed in the previous section. Our implementation directly reflects the optimizations discussed above.

14 Experimental Evaluation

A fair comparison between different families of analyses should compare both precision and speed. We have implemented a version k-CFA for a subset of R5RS Scheme and instrumented it with a possibility to optionally enable pushdown analysis, abstract garbage collection or both. Our implementation source (in Scala) and benchmarks are available:

`http://github.com/ilyasergey/reachability`

In the experiments, we have focused on the version of k-CFA with a per-state store (*i.e.*, *without* widening), as in the presence of single-threaded store, the effect of abstract GC is neutralized due to merging. For non-widened versions of k-CFA, as expected, the fused analysis does at least as well as the best of either analysis alone in terms of singleton flow sets (a good metric for program optimizability) and better than both in some cases. Also worthy of note is the dramatic reduction in the size of the abstract transition graph for the fused analysis—even on top of the already large reductions achieved by abstract garbage collection and pushdown flow analysis individually. The size of the abstract transition graph is a good heuristic measure of the temporal reasoning ability of the analysis, *e.g.*, its ability to support model-checking of safety and liveness properties (Might et al. 2007).

Program	#e	#v	k	k-CFA			k-PDCFA			k-CFA + GC			k-PDCFA + GC		
mj09	19	8	0	83	107	4	38	38	4	36	39	4	33	32	4
			1	454	812	1	44	48	1	34	35	1	32	31	1
eta	21	13	0	63	74	4	32	32	6	28	27	8	28	27	8
			1	33	33	8	32	31	8	28	27	8	28	27	8
kcfa2	20	10	0	194	236	3	36	35	4	35	34	4	35	34	4
			1	970	1935	1	87	144	2	35	34	2	35	34	2
kcfa3	25	13	0	272	327	4	58	63	5	53	52	5	53	52	5
			1	> 32662	> 88548	–	1761	4046	2	53	52	2	53	52	2
blur	40	20	0	4686	7606	4	115	146	4	90	95	10	68	76	10
			1	123	149	10	94	101	10	76	82	10	75	81	10
loop2	41	14	0	149	163	7	69	73	7	43	46	7	34	35	7
			1	> 10867	> 16040	–	411	525	3	151	163	3	145	156	3
sat	51	23	0	3844	5547	4	545	773	4	1137	1543	4	254	317	4
			1	> 28432	> 37391	–	12828	16846	4	958	1314	5	71	73	10

Table 1: Benchmark results for toy programs. The first three columns provide the name of a benchmark, the number of expressions and variables in the program in the ANF, respectively. For each of eight combinations of pushdown analysis, $k \in \{0,1\}$ and garbage collection on or off, the first two columns in a group show the number of *control states* and transitions/CRPDS edges computed during the analysis (for both less is better). The third column presents the amount of *singleton* variables, i.e, how many variables have a single lambda flow to them (more is better). Inequalities for some results of the plain k-CFA denote the case when the analysis explored more than 10^5 *configurations* (*i.e.*, control states coupled with continuations) or did not finish within 30 minutes. For such cases we do not report on singleton variables.

14.1 Plain k-CFA vs. pushdown k-CFA

In order to exercise both well-known and newly-presented instances of CESK-based CFAs, we took a series of *small* benchmarks exhibiting archetypal control-flow patterns (see Table 1). Most benchmarks are taken from the CFA literature: `mj09` is a running example from the work of Midtgaard and Jensen designed to exhibit a non-trivial return-flow behavior (Midtgaard 2007), `eta` and `blur` test common functional idioms, mixing closures and eta-expansion, `kcfa2` and `kcfa3` are two worst-case examples extracted from the proof of k-CFA's EXPTIME hardness (Van Horn and Mairson 2008), `loop2` is an example from Might's dissertation that was used to demonstrate the impact of abstract GC (Might 2007, Section 13.3), `sat` is a brute-force SAT-solver with backtracking.

14.1.1 Comparing precision

In terms of precision, the fusion of pushdown analysis and abstract garbage collection substantially cuts abstract transition graph sizes over one technique alone.

We also measure singleton flow sets as a heuristic metric for precision. Singleton flow sets are a necessary precursor to optimizations such as flow-driven inlining, type-check elimination and constant propagation. It is essential to notice that for the experiments in

Table 1 our implementation computed the sets of *values* (*i.e.*, closures) assigned to each variable (as opposed to mere *syntactic lambdas*). This is why in some cases the results computed by 0CFA appear to be better than those by 1CFA: the later one may examine *more* states with different environments, which results in exploring more different values, whereas the former one will just collapse all these values to a single lambda (Might et al. 2010). What is important is that for a fixed k the fused analysis prevails as the best-of- or better-than-both-worlds.

Running on the benchmarks, we have re-validated hypotheses about the improvements to precision granted by both pushdown analysis (Vardoulakis and Shivers 2010) and abstract garbage collection (Might 2007). Table 1 contains our detailed results on the precision of the analysis. In order to make the comparison fair, in the table we report on the numbers of *control states*, which do not contain a stack component and are the nodes of the constructed CRPDS. In the case of plain k-CFA, control states are coupled with stack pointers to obtain *configurations*, whose resulting number is significantly bigger.

The SAT-solving benchmark showed a dramatic improvement with the addition of context-sensitivity. Evaluation of the results showed that context-sensitivity provided enough fuel to eliminate most of the non-determinism from the analysis.

14.1.2 Comparing speed

In the original work on CFA2, Vardoulakis and Shivers present experimental results with a remark that the running time of the analysis is proportional to the size of the reachable states (Vardoulakis and Shivers 2010, Section 6). There is a similar correlation in our fused analysis, but with higher variance due to the live address set computation GC performs. Since most of the programs from our toy suite run for less than a second, we do not report on the absolute time. Instead, the histogram on Figure 10 presents normalized relative times of analyses' executions. To our observation the pure machine-style k-CFA is always significantly worse in terms of execution time than either with GC or push-down system, so we excluded the plain, non-optimized k-CFA from the comparison.

Our earlier implementation of a garbage-collecting pushdown analysis (Earl et al. 2012) did not fully exploit the opportunities for caching ε-predecessors, as described in Section 13. This led to significant inefficiencies of the garbage-collecting analyzer with respect to the regular k-CFA, even though the former one observed a smaller amount of states and in some cases found larger amounts of singleton variables. After this issue has been fixed, it became clearly visible that in all cases the GC-optimized analyzer is strictly faster than its non-optimized pushdown counterpart.

Although caching of ε-predecessors and ε-summary edges is relatively cheap, it is not free, since maintaining the caches requires some routine machinery at each iteration of the analyzer. This explains the loss in performance of the garbage-collecting pushdown analysis with respect to the GC-optimized k-CFA.

As it follows from the plot, fused analysis is always faster than the non-garbage-collecting pushdown analysis, and about a fifth of the time, it beats k-CFA with garbage collection in terms of performance. When the fused analysis is slower than just a GC-optimized one, it is generally not much worse than twice as slow as the next slowest analysis. Given the already

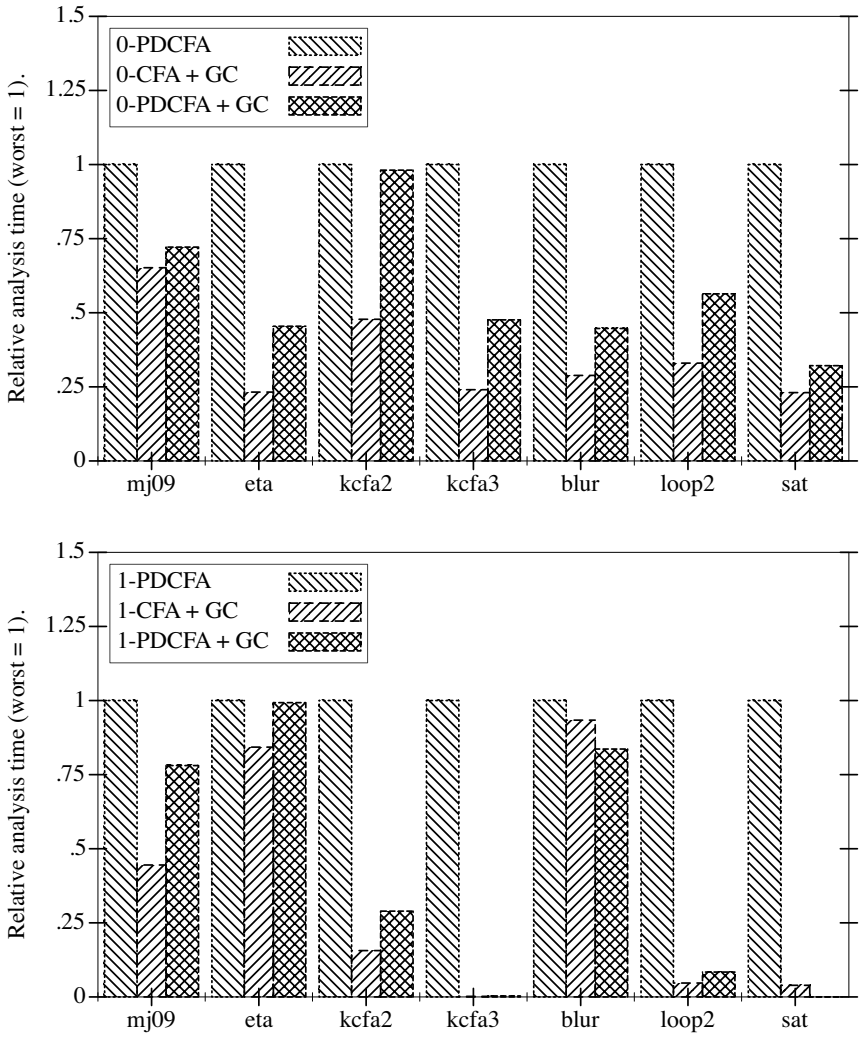

Fig. 10: Analysis times relative to worst (= 1) in class; smaller is better. At the top is the monovariant 0CFA class of analyses, at the bottom is the polyvariant 1CFA class of analyses. (Non-GC *k*-CFA omitted.)

substantial reductions in analysis times provided by collection and pushdown analysis, the amortized penalty is a small and acceptable price to pay for improvements to precision.

14.2 Analyzing real-world programs with garbage-collecting pushdown CFA

Even though our prototype implementation is just a proof of concept, we evaluated it not on a suite of toy programs, tailored for particular functional programming patterns, but on a

Program	#e	#v	!#v	$k=0$, GC off			$k=0$, GC on			$k=1$, GC off			$k=1$, GC on		
`primtest`	155	44	16	790	955	$14''$	113	127	$1''$	>43146	>54679	∞	442	562	$13''$
`rsa`	211	93	36	1267	1507	$23''$	355	407	$6''$	20746	28895	$21'$	926	1166	$28''$
`regex`	344	150	44	943	956	$54''$	578	589	$45''$	1153	1179	$88''$	578	589	$50''$
`scm2java`	1135	460	63	376	375	$13''$	376	375	$13''$	376	375	$14''$	376	375	$13''$

Table 2: Benchmark results of PDCFA on real-world programs. The first four columns provide the name of a program, the number of expressions and variables in the program in the ANF, and the number of singleton variables revealed by the analysis (same in all cases). For each of four combinations of $k \in \{0,1\}$ and garbage collection on or off, the first two columns in a group show the number of visited control states and edges, respectively, and the third one shows absolute time of running the analysis (for both less is better). The results of the analyses are presented in minutes ($'$) or seconds ($''$), where ∞ stands for an analysis, which has been interrupted due to the an execution time greater than 30 minutes.

set of real-world programs. In order to set this experiment, we have chosen four programs, dealing with numeric and symbolic computations:

- `primtest` – an implementation of the probabilistic Fermat and Solovay-Strassen primality testing in Scheme for the purpose of large prime generation;
- `rsa` – an implementation of the RSA public-key cryptosystem;
- `regex` – an implementation of a regular expression matcher in Scheme using Brzozovski derivatives (Might et al. 2011; Owens et al. 2009);
- `scm2java` – scm2java is a Scheme to Java compiler;

We ran our benchmark suite on a 2.7 GHz Intel Core i7 OS X machine with 8 Gb RAM. Unfortunately, k-CFA without global stores timed out on most of these examples (*i.e.*, did not finish within 30 minutes), so we had to exclude it from the comparison and focus on the effect of a pushdown analyzer only. Table 2 presents the results of running the benchmarks for $k \in \{0,1\}$ with a garbage collector on and off. Surprisingly, for each of the six programs, those cases, which terminated within 30 minutes, found the same number of global singleton variables.[6] However, the numbers of observed states and runtimes are indeed different in most of the cases except `scm2java`, for which all the four versions of the analysis were precise enough to actually evaluate the program: happily, there was no reuse of abstract addresses, which resulted in the absence of forking during the CRPDS construction. In other words, the program `scm2java`, which used no scalar data but strings being concatenated and was given a simple input, has been evaluated *precisely*, which is confirmed by the number of visited control states and edges.

Time-wise, the results of the experiment demonstrate the general positive effect of the abstract garbage collection in a pushdown setting, which might improve the analysis performance by more than two orders of magnitude.

[6] Of course, the numbers of states explored for each program by different analyses were different, and there were variations in *function parameters* cardinalities, which we do not report on here.

15 Discussion: Applications

Pushdown control-flow analysis offers more precise control-flow analysis results than the classical finite-state CFAs. Consequently, introspective pushdown control-flow analysis improves flow-driven optimizations (*e.g.*, constant propagation, global register allocation, inlining (Shivers 1991)) by eliminating more of the false positives that block their application.

The more compelling applications of pushdown control-flow analysis are those which are difficult to drive with classical control-flow analysis. Perhaps not surprisingly, the best examples of such analyses are escape analysis and interprocedural dependence analysis. Both of these analyses are limited by a static analyzer's ability to reason about the stack, the core competency of introspective pushdown control-flow analysis. (We leave an in-depth formulation and study of these analyses to future work.)

15.1 Escape analysis

In escape analysis, the objective is to determine whether a heap-allocated object is safely convertible into a stack-allocated object. In other words, the compiler is trying to figure out whether the frame in which an object is allocated outlasts the object itself. In higher-order languages, closures are candidates for escape analysis.

Determining whether all closures over a particular λ-term *lam* may be heap-allocated is straightforward: find the control states in the compact rooted pushdown system in which closures over *lam* are being created, then find all control states reachable from these states over only ε-edge and push-edge transitions. Call this set of control states the "safe" set. Now find all control states which are invoking a closure over *lam*. If any of these control states lies outside of the safe set, then stack-allocation may not be safe; if, however, all invocations lie within the safe set, then stack-allocation of the closure is safe.

15.2 Interprocedural dependence analysis

In interprocedural dependence analysis, the goal is to determine, for each λ-term, the set of resources which it may read or write when it is called. Might and Prabhu (2009) showed that if one has knowledge of the program stack, then one can uncover interprocedural dependencies. We can adapt that technique to work with compact rooted pushdown systems. For each control state, find the set of reachable control states along only ε-edges and pop-edges. The frames on the pop-edges determine the frames which could have been on the stack when in the control state. The frames that are live on the stack determine the procedures that are live on the stack. Every procedure that is live on the stack has a read-dependence on any resource being read in the control state, while every procedure that is live on the stack also has a write-dependence on any resource being written in the control state. In control-flow terms, this translates to "if f calls g and g accesses a, then f also accesses a."

16 Related Work

The Scheme workshop presentation of PDCFA (Earl et al. 2010) is not archival, nor were there rigorous proofs of correctness. The complete development of pushdown analysis from first principles stands as a new contribution, and it constitutes an alternative to CFA2. It goes beyond work on CFA2 by specifying specific mechanisms for reducing the complexity to polynomial time ($\mathcal{O}(n^6)$) as well. Vardoulakis (2012) sketches an approach to regain polynomial time in his dissertation, but does not give a precise bound. An immediate advantage of the complete development is its exposure of parameters for controlling polyvariance and context-sensitivity. An earlier version of this work appeared in ICFP 2012 (Earl et al. 2012). We also provide a reference implementation of control-state reachability in Haskell. We felt this was necessary to shine a light on the "dark corners" in the formalism, and in fact, it helped expose both bugs and implicit design decisions that were reflected in the revamped text of this work. The development of introspective pushdown systems is also more complete and more rigorous. We expose the critical regularity constraint absent from the ICFP 2012 work, and we specify the implementation of control-state reachability and feasible paths for conditional pushdown systems in greater detail. More importantly, this work uses additional techniques to improve the performance of the implementation and discusses those changes.

Garbage-collecting pushdown control-flow analysis draws on work in higher-order control-flow analysis (Shivers 1991), abstract machines (Felleisen and Friedman 1987) and abstract interpretation (Cousot and Cousot 1977).

Context-free analysis of higher-order programs The motivating work for our own is Vardoulakis and Shivers recent discovery of CFA2. CFA2 is a table-driven summarization algorithm that exploits the balanced nature of calls and returns to improve return-flow precision in a control-flow analysis. Though CFA2 exploits context-free languages, context-free languages are not explicit in its formulation in the same way that pushdown systems are explicit in our presentation of pushdown flow analysis. With respect to CFA2, our pushdown flow analysis is also polyvariant/context-sensitive (whereas CFA2 is monovariant/context-insensitive), and it covers direct-style.

On the other hand, CFA2 distinguishes stack-allocated and store-allocated variable bindings, whereas our formulation of pushdown control-flow analysis does not: it allocates all bindings in the store. If CFA2 determines a binding can be allocated on the stack, that binding will enjoy added precision during the analysis and is not subject to merging like store-allocated bindings. While we could incorporate such a feature in our formulation, it is not necessary for achieving "pushdownness," and in fact, it could be added to classical finite-state CFAs as well.

CFA2 has a follow-up that sacrifices its complete abstraction with the machine that only abstracts bindings in order to handle first-class control (Vardoulakis and Shivers 2011). We do not have an analogous construction since loss of complete abstraction was an anti-goal of this work. We leave an in-depth study of generalizations of CFA2's method to introspection, polyvariance and other control operators to future work.

Calculation approach to abstract interpretation Midtgaard and Jensen (2009) systematically calculate 0CFA using the Cousot-style calculational approach to abstract interpretation (Cousot 1999) applied to an ANF λ-calculus. Like the present work, Midtgaard and Jensen start with the CESK machine of Flanagan et al. (1993) and employ a reachable-states model.

The analysis is then constructed by composing well-known Galois connections to reveal a 0CFA incorporating reachability. The abstract semantics approximate the control stack component of the machine by its top element. The authors remark monomorphism materializes in two mappings: one "mapping all bindings to the same variable," the other "merging all calling contexts of the same function." Essentially, the pushdown 0CFA of Section 4 corresponds to Midtgaard and Jensen's analysis when the latter mapping is omitted and the stack component of the machine is not abstracted. However, not abstracting the stack requires non-trivial mechanisms to compute the compaction of the pushdown system.

CFL- and pushdown-reachability techniques This work also draws on CFL- and pushdown-reachability analysis (Bouajjani et al. 1997; Kodumal and Aiken 2004; Reps 1998; Reps et al. 2005). For instance, ε-closure graphs, or equivalent variants thereof, appear in many context-free-language and pushdown reachability algorithms. For our analysis, we implicitly invoked these methods as subroutines. When we found these algorithms lacking (as with their enumeration of control states), we developed rooted pushdown system compaction.

CFL-reachability techniques have also been used to compute classical finite-state abstraction CFAs (Melski and Reps 2000) and type-based polymorphic control-flow analysis (Rehof and Fähndrich 2001). These analyses should not be confused with pushdown control-flow analysis, which is computing a fundamentally more precise kind of CFA. Moreover, Rehof and Fähndrich's method is cubic in the size of the *typed* program, but the types may be exponential in the size of the program. Finally, our technique is not restricted to typed programs.

Model-checking pushdown systems with checkpoints A pushdown system with checkpoints has designated finite automata for state/frame pairs. If in a given state/frame configuration, and the automaton accepts the current stack, then execution continues. This model was first created in Esparza et al. (2003) and describes its applications to model-checking programs that use Java's `AccessController` class, and performing better data-flow analysis of Lisp programs with dynamic scope, though the specific applications are not fully explored. The algorithm described in the paper is similar to ours, but not "on-the-fly," however, so such applications would be difficult to realize with their methods. The algorithm discussed has multiple loops that enumerate all transitions within the pushdown system considered. Again this is a non-starter for higher-order languages, since up-front enumeration would conservatively suggest that any binding called would resolve to any possible function. This strategy is a sure-fire way to destroy precision and performance.

Meet-over-all-paths for conditional weighted pushdown systems A conditional pushdown system is essentially a pushdown system in which every state/frame pair is a check-

point. The two are easily interchangeable, but weighted conditional pushdown systems assign weights to reduction rules from a bounded idempotent semiring in the same manner as Reps et al. (2005). The work that introduces CWPDSs uses them for points-to analysis for Java. They solve the meet-over-all-paths problem by an incrementally translating a skeleton CFG into a WPDS and using WPDS++ (Lal and Reps 2006) to discover more points-to information to fill in call/return edges. The translation involves a heavy encoding and is not obviously correct. The killer for its use for GC is that it involves building the product automaton of all the (minimized) condition automata for the system, and interleaving the system states with the automaton's states — there are exponentially many such machines in our case, and even though the overall solution is incremental, this large automaton is pre-built. It is not obvious how to incrementalize the whole construction, nor is it obvious that the precision and performance are not negatively impacted by the repeated invocation of the WPDS solver (as opposed to a work-set solution that only considers recently changed states).

The approach to incremental solving using first-order tools is an interesting approach that we had not considered. Perhaps first-order and higher-order methods are not too far removed. It is possible that these frameworks could be extended to request transitions — or even further, checkpoint machines — on demand in order to better support higher-order languages. As we saw in this article, however, we needed access to internal data structures to compute root sets of addresses, and the ability to update a cache of such sets in these structures. The marriage could be rocky, but worth exploring in order to unite the two communities and share technologies.

Model-checking higher-order recursion schemes There is terminology overlap with work by Kobayashi (2009) on model-checking higher-order programs with higher-order recursion schemes, which are a generalization of context-free grammars in which productions can take higher-order arguments, so that an order-0 scheme is a context-free grammar. Kobyashi exploits a result by Ong (2006) which shows that model-checking these recursion schemes is decidable (but ELEMENTARY-complete) by transforming higher-order programs into higher-order recursion schemes.

Given the generality of model-checking, Kobayashi's technique may be considered an alternate paradigm for the analysis of higher-order programs. For the case of order-0, both Kobayashi's technique and our own involve context-free languages, though ours is for control-flow analysis and his is for model-checking with respect to a temporal logic. After these surface similarities, the techniques diverge. In particular, higher-order recursions schemes are limited to model-checking programs in the simply-typed lambda-calculus with recursion.

17 Conclusion

Our motivation was to further probe the limits of decidability for pushdown flow analysis of higher-order programs by enriching it with abstract garbage collection. We found that abstract garbage collection broke the pushdown model, but not irreparably so. By casting abstract garbage collection in terms of an introspective pushdown system and synthesizing

a new control-state reachability algorithm, we have demonstrated the decidability of fusing two powerful analytic techniques.

As a byproduct of our formulation, it was also easy to demonstrate how polyvariant/context-sensitive flow analyses generalize to a pushdown formulation, and we lifted the need to transform to continuation-passing style in order to perform pushdown analysis.

Our empirical evaluation is highly encouraging: it shows that the fused analysis provides further large reductions in the size of the abstract transition graph—a key metric for interprocedural control-flow precision. And, in terms of singleton flow sets—a heuristic metric for optimizability—the fused analysis proves to be a "better-than-both-worlds" combination.

Thus, we provide a sound, precise and polyvariant introspective pushdown analysis for higher-order programs.

Acknowledgments

We thank the anonymous reviewers of ICFP 2012 and JFP for their detailed reviews, which helped to improve the presentation and technical content of the paper. Tim Smith was especially helpful with his knowledge of stack automata. This material is based on research sponsored by DARPA under the programs Automated Program Analysis for Cybersecurity (FA8750-12-2-0106) and Clean-Slate Resilient Adaptive Hosts (CRASH). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

References

Bouajjani, A., J. Esparza, and O. Maler (1997). Reachability analysis of pushdown automata: Application to Model-Checking. In *Proceedings of the 8th International Conference on Concurrency Theory*, CONCUR '97, pp. 135–150. Springer-Verlag.

Cousot, P. (1999). The calculational design of a generic abstract interpreter. In M. Broy and R. Steinbrüggen (Eds.), *Calculational System Design*. NATO ASI Series F. IOS Press, Amsterdam.

Cousot, P. and R. Cousot (1977). Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages*, pp. 238–252. ACM Press.

Earl, C., M. Might, and D. Van Horn (2010). Pushdown Control-Flow analysis of Higher-Order programs. In *Workshop on Scheme and Functional Programming*.

Earl, C., I. Sergey, M. Might, and D. Van Horn (2012). Introspective pushdown analysis of higher-order programs. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming (ICFP 2012)*, ICFP '12, pp. 177–188. ACM.

Esparza, J., A. Kucera, and S. Schwoon (2003). Model checking LTL with regular valuations for pushdown systems. *Inf. Comput. 186*(2), 355–376.

Felleisen, M. and D. P. Friedman (1987). A calculus for assignments in higher-order languages. In *POPL '87: Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pp. 314+. ACM.

Flanagan, C., A. Sabry, B. F. Duba, and M. Felleisen (1993, June). The essence of compiling with continuations. In *PLDI '93: Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, pp. 237–247. ACM.

Ginsburg, S., S. A. Greibach, and M. A. Harrison (1967). One-way stack automata. *Journal of the ACM 14*(2), 389–418.

Johnson, J. I. and D. Van Horn (2013). Concrete semantics for pushdown analysis: The essence of summarization. In *HOPA 2013: Workshop on higher-order program analysis*.

Kobayashi, N. (2009, January). Types and higher-order recursion schemes for verification of higher-order programs. *SIGPLAN Not. 44*(1), 416–428.

Kodumal, J. and A. Aiken (2004, June). The set constraint/CFL reachability connection in practice. In *PLDI '04: Proceedings of the ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation*, pp. 207–218. ACM.

Lal, A. and T. W. Reps (2006). Improving pushdown system model checking. In T. Ball and R. B. Jones (Eds.), *CAV*, Volume 4144 of *Lecture Notes in Computer Science*, pp. 343–357. Springer.

Li, X. and M. Ogawa (2010). Conditional weighted pushdown systems and applications. In J. P. Gallagher and J. Voigtländer (Eds.), *PEPM*, pp. 141–150. ACM.

Melski, D. and T. W. Reps (2000, October). Interconvertibility of a class of set constraints and context-free-language reachability. *Theoretical Computer Science 248*(1-2), 29–98.

Midtgaard, J. (2007). *Transformation, Analysis, and Interpretation of Higher-Order Procedural Programs*. Ph. D. thesis, University of Aarhus.

Midtgaard, J. and T. P. Jensen (2009). Control-flow analysis of function calls and returns by abstract interpretation. In *ICFP '09: Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming*, pp. 287–298. ACM.

Might, M. (2007, June). *Environment Analysis of Higher-Order Languages*. Ph. D. thesis, Georgia Institute of Technology.

Might, M., B. Chambers, and O. Shivers (2007, January). Model checking via Gamma-CFA. In *Verification, Model Checking, and Abstract Interpretation*, pp. 59–73.

Might, M., D. Darais, and D. Spiewak (2011). Parsing with derivatives: a functional pearl. In *ICFP '11: Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming*, pp. 189–195. ACM.

Might, M. and P. Manolios (2009). A posteriori soundness for non-deterministic abstract interpretations. In *Proceedings of the 10th International Conference on Verification, Model Checking, and Abstract Interpretation*, VMCAI '09, pp. 260–274. Springer-Verlag.

Might, M. and T. Prabhu (2009). Interprocedural dependence analysis of higher-order programs via stack reachability. In *Proceedings of the 2009 Workshop on Scheme and Functional Programming*.

Might, M. and O. Shivers (2006a). Environment analysis via Delta-CFA. In *Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2006)*, pp. 127–140. ACM.

Might, M. and O. Shivers (2006b). Improving flow analyses via Gamma-CFA: Abstract garbage collection and counting. In *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming (ICFP 2006)*, pp. 13–25. ACM.

Might, M., Y. Smaragdakis, and D. Van Horn (2010). Resolving and exploiting the k-CFA paradox: Illuminating functional vs. object-oriented program analysis. In *PLDI '10: Proceedings of the 2010 ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '10, pp. 305–315. ACM Press.

Ong, C. H. L. (2006). On Model-Checking trees generated by Higher-Order recursion schemes. In *21st Annual IEEE Symposium on Logic in Computer Science (LICS'06)*, LICS, pp. 81–90. IEEE.

Owens, S., J. Reppy, and A. Turon (2009). Regular-expression derivatives re-examined. *Journal of Functional Programming 19*(02), 173–190.

Rehof, J. and M. Fähndrich (2001). Type-based flow analysis: From polymorphic subtyping to CFL-reachability. In *POPL '01: Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 54–66. ACM.

Reps, T. (1998, December). Program analysis via graph reachability. *Information and Software Technology 40*(11-12), 701–726.

Reps, T., S. Schwoon, S. Jha, and D. Melski (2005, October). Weighted pushdown systems and their application to interprocedural dataflow analysis. *Science of Computer Programming 58*(1-2), 206–263.

Rounds, W. C. (1973). Complexity of recognition in intermediate level languages. In *Switching and Automata Theory, 1973. SWAT '08. IEEE Conference Record of 14th Annual Symposium on*, pp. 145–158.

Shivers, O. G. (1991). *Control-Flow Analysis of Higher-Order Languages*. Ph. D. thesis, Carnegie Mellon University.

Sipser, M. (2005, February). *Introduction to the Theory of Computation* (2 ed.). Cengage Learning.

Van Horn, D. and H. G. Mairson (2008). Deciding *k*CFA is complete for EXPTIME. In *ICFP '08: Proceeding of the 13th ACM SIGPLAN International Conference on Functional Programming*, pp. 275–282. ACM.

Van Horn, D. and M. Might (2012). Systematic abstraction of abstract machines. *Journal of Functional Programming 22*(Special Issue 4-5), 705–746.

Vardoulakis, D. (2012). *CFA2: Pushdown Flow Analysis for Higher-Order Languages*. Ph. D. thesis, Northeastern University.

Vardoulakis, D. and O. Shivers (2010). CFA2: A Context-Free approach to Control-Flow analysis. In A. D. Gordon (Ed.), *Programming Languages and Systems*, Volume 6012 of *Lecture Notes in Computer Science*, Chapter 30, pp. 570–589. Springer Berlin Heidelberg.

Vardoulakis, D. and O. Shivers (2011). Pushdown flow analysis of first-class control. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming*, ICFP '11, pp. 69–80. ACM.

Wright, A. K. and S. Jagannathan (1998, January). Polymorphic splitting: An effective polyvariant flow analysis. *ACM Transactions on Programming Languages and Systems 20*(1), 166–207.

18 Full Proofs

18.1 Pushdown reachability

Proof of Theorem 8.1. The space *ICRPDS* is further constrained than stated in the main article:

$$ICRPDS = \left\{ ((S,E),H,(\Delta S,\Delta E,\Delta H)) : \begin{array}{l} \bigcup\{\{s,s'\} : s \rightarrowtail s' \in H\} \subseteq S, \\ \Delta S \cap S = \emptyset, \Delta E \cap E = \emptyset, \text{ and } \Delta H \cap H = \emptyset \end{array} \right\}$$

For this section we assume

$$M = (Q,\Gamma,\delta,q_0) \in \mathbb{RPDS}$$
$$G = ((S,E),H,(\Delta S,\Delta E),\Delta H) \in \mathbb{CRPDS} \text{ where } (S,E) \subseteq (Q,\delta)$$
$$\text{and } q_0 = s_0$$

Let $\mathbb{ORD}$ be the class of ordinal numbers. We define a termination measure on the fixed-point computation of $\mathscr{F}'((Q,\Gamma,_,_))$, $d : \mathbb{CRPDS} \to \mathbb{ORD}$.

$$d((S,E),H,(\Delta S,\Delta E,\Delta H)) = (2^{|Q|^2 \cdot |\Gamma|} - |E|)\omega + (2^{|Q|^2} - |H|)$$

Lemma 18.1 (*Termination*)
Either $G = \mathscr{F}'(M)(G)$ or $d(\mathscr{F}'(M)(G)) \prec d(G)$

Proof
If both ΔE and ΔH are empty, there are no additions made to S, E or H, meaning G is a fixed point. Otherwise, due to the non-overlap condition, one or both of E and H grow, meaning the ordinal is smaller. □

A corollary is that the fixed-point has empty ΔE and ΔH.

Lemma 18.2 (*Key lemma for PDS reachability*)
If $inv(G)$ then $inv(\mathscr{F}'(M)(G))$

Proof
All additional states and edges come from ΔE_i and ΔH_i for $i \in [0..4]$, so by cases on the sources of edges:

Case $s \overset{g}{\rightarrowtail} s' \in \Delta E_0, s'' \rightarrowtail s''' \in \Delta H_0$**.**
By definition of *sprout* and path extension.

Case $s \overset{g}{\rightarrowtail} s' \in \Delta E_1, s'' \rightarrowtail s''' \in \Delta H_1$**.**

If $g \equiv \hat{\phi}_-$, then by definition of *addPush* there are $q \overset{\hat{\phi}_+}{\rightarrowtail} q' \in \Delta E$, $q' \rightarrowtail s \in H$, such that $(s,\hat{\phi}_+,s') \in \delta$.
Let $\vec{g}$ be the witness of the invariant on $q' \rightarrowtail s$ given from definition of *inv*. Let $\hat{\kappa}$ be arbitrary. We have $[\hat{\phi}_+\vec{g}\hat{\phi}_-] = \varepsilon$. We also have $(q,\hat{\kappa}) \overset{\hat{\phi}_+\vec{g}\hat{\phi}_-}{\underset{M}{\longmapsto}}{}^* (s',\hat{\kappa})$. Root reachability follows from path concatenation with the root path from $(q,\hat{\kappa}) \overset{\hat{\phi}_+}{\underset{M}{\longmapsto}} (q',\hat{\phi}\hat{\kappa})$ from *inv*.
The balanced path for $s'' \rightarrowtail s'''$ comes from a similar push edge from ΔE and concatenation with the path from the invariant on H.

Case $s'' \rightarrowtail s''' \in \Delta H_2$**.**

By definition of *addPop*, $\Delta E_2 = \emptyset$ and there are $q \overset{\hat{\phi}_-}{\rightarrowtail} s''' \in \Delta E$, $q' \rightarrowtail q \in H$ such that $s'' \overset{\hat{\phi}_+}{\rightarrowtail} q' \in E$. Let $\vec{g}$ be the witness of the invariant on $q' \rightarrowtail q$. Let $\hat{\kappa}$ be arbitrary. We know by the invariant on E, $(s'', \hat{\kappa}) \underset{M}{\overset{\hat{\phi}_+ \vec{g} \hat{\phi}_-}{\longmapsto}}{}^{*} (s''', \hat{\kappa})$ and $[\hat{\phi}_+ \vec{g} \hat{\phi}_-] = \varepsilon$.

Case $s \overset{g}{\rightarrowtail} s' \in \Delta E_3 \cup \Delta E_4, s'' \rightarrowtail s''' \in \Delta H_3 \cup \Delta H_4$**.**

Follows from definition of *inv* and path concatenation, following similar reasoning as above cases.

□

We define "π is a subtrace of π'," $\pi \sqsubseteq \pi'$

$$\frac{}{(s', \hat{\kappa}') \underset{M}{\overset{\langle\rangle}{\longmapsto}}{}^{*} (s', \hat{\kappa}') \sqsubseteq (s, \hat{\kappa}) \underset{M}{\overset{\vec{g}}{\longmapsto}}{}^{*} (s', \hat{\kappa}')} \qquad \frac{\pi \sqsubseteq (s, \hat{\kappa}) \underset{M}{\overset{\vec{g}}{\longmapsto}}{}^{*} (s', \hat{\kappa}') \quad (s', g', s'') \in \delta}{\pi \sqsubseteq (s, \hat{\kappa}) \underset{M}{\overset{\vec{g}}{\longmapsto}}{}^{*} (s', \hat{\kappa}') \underset{M}{\overset{g'}{\longmapsto}} (s'', [\hat{\kappa}'_+ g'])}$$

$$\frac{(s, \hat{\kappa}) \underset{M}{\overset{\vec{g}}{\longmapsto}}{}^{*} (s', \hat{\kappa}') \sqsubseteq (s''', \hat{\kappa}'') \underset{M}{\overset{\vec{g''}}{\longmapsto}}{}^{*} (s', \hat{\kappa}') \quad (s', g', s'') \in \delta}{(s, \hat{\kappa}) \underset{M}{\overset{\vec{g}}{\longmapsto}}{}^{*} (s', \hat{\kappa}') \underset{M}{\overset{g'}{\longmapsto}} (s'', [\hat{\kappa}'_+ g']) \sqsubseteq (s''', \hat{\kappa}'') \underset{M}{\overset{\vec{g''}}{\longmapsto}}{}^{*} (s', \hat{\kappa}') \underset{M}{\overset{g'}{\longmapsto}} (s'', [\hat{\kappa}'_+ g'])}$$

Theorem 8.1 is a corollary of the following theorem.

Theorem 18.1
$\mathrm{lfp}(\mathscr{F}'(M)) = (\mathscr{C}(M), \mathscr{ECG}(M), (\emptyset, \emptyset), \emptyset)$

Proof
($\subseteq$): Directly from 18.2.
($\supseteq$): Let $\pi \equiv (s_0, \langle\rangle) \underset{M}{\overset{\vec{g}}{\longmapsto}}{}^{*} (s, \hat{\kappa})$ be an arbitrary path in $\mathscr{C}(M)$ (the inclusion of root is not a restriction due to the definition of CRPDSs). Let $n \in Nats$ be such that $\mathrm{lfp}(\mathscr{F}'(M)) = \mathscr{F}'(M)^n$. We show

- the same path through G,
- for each $s \in S$, $s \overset{g}{\rightarrowtail} s' \in E$, $s \rightarrowtail s' \in H$, there is an $m < n$ such that $s \in \Delta S_m$ $s \overset{g}{\rightarrowtail} s' \in \Delta E_m$ $s \rightarrowtail s' \in \Delta H_m$ respectively, where $\mathscr{F}'(M)^m = (G_m, H_m, (\Delta S_m, \Delta E_m, \Delta H_m))$, and
- all non-empty balanced subtraces have edges in H: $\forall (s_b, \hat{\kappa}) \underset{M}{\overset{\vec{g_\varepsilon}}{\longmapsto}}{}^{*} (s_a, \hat{\kappa}) \sqsubseteq \pi.\vec{g}_\varepsilon \neq \langle\rangle \wedge [\vec{g}_\varepsilon] = \varepsilon \implies s_b \rightarrowtail s_a \in H$.

By induction on π,

Case Base: s_0**.**

Follows by definition of $\mathscr{F}'$. No non-empty balanced subtrace.

Case Induction step: $(s_0, \langle\rangle) \underset{M}{\overset{\vec{g'}}{\longmapsto}}{}^{*} (s', \hat{\kappa}) \underset{M}{\overset{g''}{\longmapsto}} (s, [\hat{\kappa}_+ g''])$**.**

By IH, $(s_0, \langle\rangle) \underset{G}{\overset{\vec{g'}}{\longmapsto}}{}^{*} (s', \hat{\kappa})$. By cases on g'':

Case γ_+.

Let m be the witness for s' by the IH. By definition of $\mathscr{F}'$, $(s', \hat{\kappa}) \overset{g''}{\underset{M}{\longmapsto}} (s, [\hat{\kappa}_+ g''])$ is in ΔE_{m+1} and E_{m+2} (and thus $s \in \Delta S_{m+1}$ and S_{m+2}). Thus the path is constructible through G_n. All balanced subtraces carry over from IH, since the last push edge cannot end a balanced path.

Case ε.

The path is constructible the same as γ_+. Let m be the witness used in the path construction. Let $\pi' \equiv (s_b, \hat{\kappa}) \overset{\vec{g_\varepsilon}}{\underset{M}{\longmapsto}}{}^* (s_e, \hat{\kappa})$ be an arbitrary non-empty balanced subtrace. If $s_e \neq s$, then the IH handles it. Otherwise, $\vec{g_\varepsilon} = \vec{g'_\varepsilon}\varepsilon$. If $s_b = s'$, then the ε-edge is added by *sprout* (so the witness number is $m+1$). If not, then there is a balanced subtrace $(s_b, \hat{\kappa}) \overset{\vec{g'_\varepsilon}}{\underset{M}{\longmapsto}}{}^* (s', \hat{\kappa})$, thus $s_b \rightarrowtail s' \in H$. Let m' be the witness for $s_b \rightarrowtail s' \in \Delta H_{m'}$. Then $s_b \rightarrowtail s \in \Delta_{\max\{m,m'\}+1}$ by definition of *addEmpty*.

Case γ_-.

Since $[\vec{g}]$ is defined, there is a push edge in the trace (call it $s_u \overset{\gamma_+}{\underset{M}{\longmapsto}} s_v$) with a (possibly empty) balanced subtrace following to s'. Thus by IH, there are some m, m' such that $s_u \overset{\gamma_+}{\rightarrowtail} s_v \in E_m$, (if the subtrace is non-empty) $s_v \rightarrowtail s' \in H_{m'}$ If $m \geq m'$ by definition of *addPush*, $s' \overset{\gamma_-}{\rightarrowtail} s \in \Delta E_{m+1}$. Otherwise, the edge is in $E_{m'}$ and by definition of *addEmpty*, $s' \overset{\gamma_-}{\rightarrowtail} s \in \Delta E_{m'+1}$.

Let $\pi' \equiv (s_b, \hat{\kappa}) \overset{\vec{g_\varepsilon}}{\underset{M}{\longmapsto}}{}^* (s_e, \hat{\kappa})$ be an arbitrary non-empty balanced subtrace. If $s_e \neq s$, the IH handles it. Otherwise, $\vec{g_\varepsilon} = \vec{g'_\varepsilon}\gamma_+\vec{g''_\varepsilon}\gamma_-$ and $\pi' \equiv (s_b, \hat{\kappa}) \overset{\vec{g'_\varepsilon}}{\underset{M}{\longmapsto}}{}^* (s_u, \hat{\kappa}) \overset{\gamma_+}{\underset{M}{\longmapsto}} (s_v, \gamma\hat{\kappa}) \overset{\vec{g'_\varepsilon}}{\underset{M}{\longmapsto}}{}^* (s', \gamma\hat{\kappa}) \overset{\gamma_-}{\underset{M}{\longmapsto}} (s, \hat{\kappa})$. $s_u \rightarrowtail s$ is added to $\Delta H_{\max\{m,m'\}+1}$ and thus $s_b \rightarrowtail s_u$ is in $H_{\max\{m,m'\}+3}$.

□

18.2 RIPDS reachability

We use metafunction $\bullet{+}{+}\bullet : \mathit{Cont} \times \mathit{Cont} \to \mathit{Cont}$ to aid proofs:

$$\varepsilon{+}{+}\hat{\kappa} = \hat{\kappa}$$
$$\phi : \hat{\kappa}{+}{+}\hat{\kappa}' = \phi : (\hat{\kappa}{+}{+}\hat{\kappa}')$$

$$\mathit{split}(\varepsilon) = [\not\varepsilon, \varepsilon]$$
$$\mathit{split}(\phi : \hat{\kappa}) = [\not\varepsilon\hat{\kappa}, \phi]$$

Lemma 18.3 (*Down spin*)
For $(q, \varepsilon, \varepsilon, \downarrow, \varepsilon, q) \in \delta$, $(q, [\hat{\kappa}_B{+}{+}\hat{\kappa}_{B'}, \hat{\kappa}_T], w) \longmapsto^* (q, [\hat{\kappa}_B, \hat{\kappa}_{B'}{+}{+}\hat{\kappa}_T], w)$

Proof
By induction on $\hat{\kappa}_{B'}$.

Case Base: $\hat{\kappa}_B = \varepsilon$.
Reflexivity.
Case Induction step: $\hat{\kappa}_{B'} = \hat{\kappa}\phi$.
By δ, $(q, [\hat{\kappa}_B{+}{+}\hat{\kappa}_{B'}, \hat{\kappa}_T], w) \longmapsto (q, [\hat{\kappa}_B{+}{+}\hat{\kappa}, \phi\,\hat{\kappa}_T], w)$. By IH, $(q, [\hat{\kappa}_B{+}{+}\hat{\kappa}, \phi\,\hat{\kappa}_T], w)$ $\longmapsto^* (q, [\hat{\kappa}_B, \hat{\kappa}{+}{+}\phi\,\hat{\kappa}_T], w)$. This final configuration is the same as $(q, [\hat{\kappa}_B, \hat{\kappa}_{B'}{+}{+}\hat{\kappa}_T], w)$.

□

Lemma 18.4 (*gadget correctness*)
For $(\delta, S) = gadget(s, \hat{K}, g, s')$, $(s, split(\hat{\kappa}), w) \underset{\delta}{\longmapsto}{}^* (s', split([\hat{\kappa}_{+}g]), w)$ iff $\hat{\kappa} \in \hat{K}$ and $[\hat{\kappa}_{+}g]$ defined.

Proof
($\Rightarrow$): By inversion on the rules for δ, the path must go through three stages: the down-spin, the middle path, and the pop-off. By 18.3, $(s, split(\hat{\kappa}), w) \longmapsto (q_{\text{down}}, [\text{¢}\,\hat{\kappa}, \$], w) \longmapsto^*$ $(q_{\text{down}}, [\varepsilon, \text{¢}\,\hat{\kappa}\$], w)$. Then the $(q_{\text{down}}, \text{¢}, \varepsilon, \uparrow, \varepsilon, q_0)$ rule must apply. We can construct an accepting path in the machine recognizing $\hat{K}$ from the middle path via the following lemma: $(q_0, [\text{¢}, \hat{\kappa}\$], w) \underset{\delta}{\overset{\hat{\kappa}'}{\longmapsto}}{}^* (q, [\text{¢}\,\hat{\kappa}', \hat{\kappa}''\$], w)$ implies $(q_0, \hat{\kappa}) \underset{N}{\overset{\hat{\kappa}'}{\longmapsto}}{}^* (q, \hat{\kappa}'')$. Proof by induction.

Then $(q, \$, \varepsilon, \cdot, \$_{-}, q_{\text{out}})$ must apply, and then $(q_{\text{out}}, \varepsilon, \varepsilon, g, s')$ must apply, meaning that $[\hat{\kappa}_{+}g]$ is defined.

($\Leftarrow$): Since $\hat{K}$ is regular, there must be a path in the chosen NFA $N = (Q, \Sigma, \delta_N, q_0, F)$ from q_0 to a final state $q \in F$, $(q_0, \hat{\kappa}) \underset{N}{\longmapsto}{}^* (q, \varepsilon)$.

In the first stage, $(s, split(\hat{\kappa}), w) \longmapsto^* (q_0, [\text{¢}, \hat{\kappa}\$], w)$.
The follows first by the $(s, \varepsilon, \varepsilon, \cdot, \$_{+}, q_{\text{down}})$ transition, then by 18.3 $(q_{\text{down}}, split(\hat{\kappa}\$), w)$ $\longmapsto^* (q_{\text{down}}, [\varepsilon, \text{¢}\,\hat{\kappa}\$], w)$, finally by the $(q_{\text{down}}, \text{¢}, \varepsilon, \uparrow, \varepsilon, q_0)$ rule.

In the second stage we construct a path $(q_0, [\text{¢}, \hat{\kappa}\$], w) \longmapsto^* (q, [\text{¢}\,\hat{\kappa}, \$], w)$, from an accepting path in N: $(q_0, \hat{\kappa}) \longmapsto^* (q, \varepsilon)$ where $q \in F$. The statement we can induct on to get this is $(q_0, \hat{\kappa}) \underset{N}{\overset{\hat{\kappa}'}{\longmapsto}}{}^* (q, \hat{\kappa}'')$ implies $(q_0, [\text{¢}, \hat{\kappa}\$], w) \underset{\delta}{\overset{\hat{\kappa}'}{\longmapsto}}{}^* (q, [\text{¢}\,\hat{\kappa}', \hat{\kappa}''\$], w)$.

Case Base: $\hat{\kappa}' = \varepsilon, q = q_0, \hat{\kappa}'' = \hat{\kappa}$.
Reflexivity.
Case Induction step: $\hat{\kappa}' = \hat{\kappa}'''\phi_\varepsilon, (q_0, \hat{\kappa}) \underset{N}{\overset{\hat{\kappa}'''}{\longmapsto}}{}^* (q', \hat{\kappa}'''') \underset{N}{\overset{\phi_\varepsilon}{\longmapsto}} (q, \hat{\kappa}'')$.
By IH, $(q_0, [\text{¢}, \hat{\kappa}\$], w) \longmapsto^* (q', [\text{¢}\,\hat{\kappa}''', \hat{\kappa}''''\$], w)$. If $\phi_\varepsilon = \varepsilon$, then $\hat{\kappa}''' = \hat{\kappa}'$, $\hat{\kappa}'''' = \hat{\kappa}''$ and we apply the $(q', \varepsilon, \varepsilon, \cdot, \varepsilon, q)$ rule to get to $(q, [\text{¢}\,\hat{\kappa}', \hat{\kappa}''\$], w)$. Otherwise, $\hat{\kappa}' = \hat{\kappa}'''\phi$ and we apply the $(q', \phi, \varepsilon, \uparrow, \varepsilon, q)$ rule to get to $(q, [\text{¢}\,\hat{\kappa}', \hat{\kappa}''\$], w)$.

In the third and final stage, $(q_{\text{out}}, [\text{¢}\,\hat{\kappa}, \$], w) \longmapsto (q_{\text{out}}, split(\hat{\kappa}), w)$ and since $[\hat{\kappa}_{+}g]$ is defined, we reach the final state by $(q_{\text{out}}, \varepsilon, \varepsilon, \cdot, g, s')$. □

Lemma 18.5 (*Checking lemma*)
If $(q, a, a, \uparrow, \varepsilon, q) \in \delta$ and $(q, [\hat{\kappa}_B, \hat{\kappa}_{T'}{+}{+}\hat{\kappa}_T\$], w) \longmapsto^* (q, [\hat{\kappa}_B{+}{+}\hat{\kappa}_{T'}, \hat{\kappa}_T\$], w')$ (through the one rule) then $w = \hat{\kappa}_{T'}w'$.

Proof

Simple induction. □

Lemma 18.6 (*Stack machine correctness*)

For all $M \in \mathbb{CPDS}$, $G \in \mathbb{CCPDS}$, $q \in G$, if $G \sqsubseteq \mathscr{CC}(M)$ then

$$\mathscr{L}(Stacks(G)(q)) = \left\{ \hat{\kappa} : (q_0, \langle\rangle) \overset{\overrightarrow{\hat{K},g}}{\underset{G}{\longmapsto}}{}^{*} (q, \hat{\kappa}) \right\}.$$

Proof

($\subseteq$): Let $(s_{\text{start}}, [\varepsilon, \varepsilon], \hat{\kappa}) \longmapsto^{*} (s_{\text{final}}, split(\hat{\kappa}), \varepsilon)$ be an accepting path for $\hat{\kappa} \in \mathscr{L}(Stacks(G)(q))$. We inductively construct a corresponding path in G that realizes $\hat{\kappa}$. We first see that the given path is split into three phases: setup, gadgetry, checking. The first step must be $(s_{\text{start}}, \varepsilon, \varepsilon, \cdot, \not\in_+, s_0)$, which we call setup. The only final state must be preceded by s_{check}, s_{down}, and the final occurrence of s, which we call checking. Thus the middle phase is a trace from s_0 to s. This must be through gadgets, which are disjoint for each rule of the IPDS, and thus each edge in G.
$(s_{\text{start}}, [\varepsilon, \varepsilon], \hat{\kappa}) \longmapsto (s_0, [\varepsilon, \not\in], \hat{\kappa}) \longmapsto^{*} (s, split(\hat{\kappa}), \hat{\kappa}) \longmapsto$
$(s_{\text{down}}, [\not\in \hat{\kappa}, \$], \hat{\kappa}) \longmapsto^{*} (s_{\text{down}}, [\varepsilon, \not\in \hat{\kappa}\$], \hat{\kappa}) \longmapsto$
$(s_{\text{check}}, [\not\in, \hat{\kappa}\$], \hat{\kappa}) \longmapsto^{*} (s_{\text{check}}, [\not\in \hat{\kappa}, \$], \varepsilon) \longmapsto (s_{\text{final}}, split(\hat{\kappa}), \varepsilon)$
We induct on the path through gadgets, s_0 to s in the above path, invoking 18.4 at each step.

($\supseteq$): Simple induction between setup and teardown, applying 18.4. □

Proof of Theorem 12.1

Proof

The finiteness of the state space and monotonicity of $\mathscr{F}$ ensures the least fixed point exists. $\text{lfp}(\mathscr{F}(M)) \subseteq \mathscr{CC}(M)$ follows from 18.6 and the definition of $\mathscr{F}$.

To prove $\mathscr{CC}(M) \subseteq \text{lfp}(\mathscr{F}(M))$, suppose not. Then there must be a path $(s_0, \langle\rangle) \overset{\overrightarrow{\hat{K},g}}{\underset{M}{\longmapsto}}{}^{*}$ $(s, \hat{\kappa}) \overset{\hat{K}',g'}{\underset{M}{\longmapsto}} (s', [\hat{\kappa}_+ g'])$ where the final edge is the first edge not in $\text{lfp}(\mathscr{F}(M))$.

By definition of $\mathscr{CC}$, $\hat{\kappa} \in \hat{K}'$ and $(s, \hat{K}', g, s') \in \delta$. Since $\hat{\kappa}$ is realizable at s in G, by definition of $\mathscr{F}$ and 18.6, $(s, \hat{\kappa}) \overset{\hat{K}',g}{\underset{G}{\longmapsto}} (s', [\hat{\kappa}_+ g])$ contra the assumption. Thus $\mathscr{CC}(M) \subseteq \text{lfp}(\mathscr{F}(M))$ holds by contradiction. □

We first prove an invariant of $\hat{f} : \mathsf{Exp} \to \widehat{System}_\Gamma \xrightarrow{\text{mon}} \widehat{System}_\Gamma$, where $\widehat{System}_\Gamma = \mathbb{ICRPDS} \times \mathscr{P}(\widehat{OPState} \times \widehat{OPState})$

$$\mathscr{I}_\Gamma(e) = ((e, \bot, \bot), \langle\rangle)$$
$$\mathscr{I}'_\Gamma(e) = (((e, \bot, \bot), \emptyset), \langle\rangle)$$
$$\langle \hat{\phi}_1, \ldots, \hat{\phi}_n \rangle_+ = \hat{\phi}_{1+} \ldots \hat{\phi}_{n+}$$
$$inv_\Gamma : \mathsf{Exp} \to \widehat{System}_\Gamma \to Prop$$
$$inv_\Gamma(e)(\overbrace{(\hat{P}, \hat{E}), \hat{H}}^{G}) = (\hat{P} = \bigcup \left\{ \{\hat{\Omega}, \hat{\Omega}'\} : \hat{\Omega} \overset{g}{\rightarrowtail} \hat{\Omega}' \in \hat{E} \right\})$$
$$\wedge \forall (\hat{\psi}, A) \overset{g}{\rightarrowtail} (\hat{\psi}', A') \in \hat{E}.\text{let } \hat{K} = \{\hat{\kappa} : StackRoot(\hat{\kappa}) = A\} \text{ in}$$
$$\forall \hat{\kappa} \in \left\{ \hat{\kappa} \in \hat{K} : [\hat{\kappa}_+ g] \text{ defined} \right\}.StackRoot([\hat{\kappa}_+ g]) = A'$$
$$\wedge (\hat{\psi}, \hat{\kappa}) \underset{M}{\overset{\hat{K}, g}{\longmapsto}} (\hat{\psi}', [\hat{\kappa}_+ g])$$
$$\wedge \forall (\hat{\psi}, _) \rightarrowtail (\hat{\psi}', _) \in \hat{H}.\exists \overrightarrow{\hat{K}, g}.[\vec{g}] = \varepsilon \wedge (\hat{\psi}, \langle\rangle) \underset{M}{\overset{\overrightarrow{\hat{K}, g}^*}{\longmapsto}} (\hat{\psi}', \langle\rangle)$$
$$\text{where } M = \widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}}'(e)$$

Lemma 18.7 ($\hat{f}$ *invariant*)
For all e, if $inv_\Gamma(e)(G)$ then $inv_\Gamma(e)(\hat{f}_e(G))$

Proof
Same structure as in Lemma 18.2 without reasoning about worklists. □

Proof of Theorem 12.2

Proof
Let $M = \widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}}'(e)$, $G = \mathscr{C}\mathscr{C}(M)$ and $G' = ((\hat{P}, \hat{E}), \hat{H}) = \text{lfp}(\hat{f}_e)$.
(G' approximates G):

We strengthen the statement to $\pi \equiv \mathscr{I}_\Gamma(e) \underset{G}{\overset{\overrightarrow{\hat{K}, g}}{\longmapsto}}{}^* (\hat{\psi}, \hat{\kappa})$ implies

- $\mathscr{I}'_\Gamma(e) \underset{G'}{\overset{\vec{g}}{\longmapsto}}{}^* ((\hat{\psi}, StackRoot(\hat{\kappa})), \hat{\kappa})$.
- for all $(\hat{\psi}, \hat{\kappa}) \underset{G}{\overset{\overrightarrow{\hat{K}, g}^*}{\longmapsto}} (\hat{\psi}', [\hat{\kappa}_+ \vec{g}]) \underset{G}{\overset{\hat{K}', g'}{\longmapsto}} (\hat{\psi}'', [\hat{\kappa}_+ \vec{g} g']) \sqsubseteq \pi$, if $[\vec{g} g'] = \varepsilon$, then $\exists \hat{\kappa} \in \hat{K}$ and $(\hat{\psi}, StackRoot(\hat{\kappa})) \rightarrowtail (\hat{\psi}'', StackRoot([\hat{\kappa}_+ \vec{g} g'])) \in \hat{H}$

By induction on π,

Case Base: $\mathscr{I}_\Gamma(e)$**.**
By definition of $\hat{f}_e$, $\mathscr{I}'_\Gamma(e) = (\hat{\Omega}_0, \langle\rangle)$, $\hat{\Omega}_0 \in \hat{P}$. First goal holds by definition of $StackRoot(\langle\rangle)$ and reflexivity. Second goal vacuously true.

Case Induction step: $((e, \bot, \bot), \langle\rangle) \underset{G}{\overset{\overrightarrow{\hat{K}', g'}}{\longmapsto}}{}^* (\hat{\psi}', [\vec{g'}]) \underset{G}{\overset{\hat{K}'', g''}{\longmapsto}} (\hat{\psi}, \hat{\kappa})$**.**

Let $A = StackRoot([\vec{g'}])$. By IH, $\mathscr{I}'_\Gamma(e) \underset{G'}{\overset{\vec{g'}}{\longmapsto}}{}^* ((\hat{\psi}', A), [\vec{g'}])$.

Let $\hat{K}_{\text{root}} = \{\hat{\kappa} : \mathit{StackRoot}(\hat{\kappa}) = A\}$ By definition of $\widehat{\mathscr{I}\mathscr{P}\mathscr{D}\mathscr{S}}'$ and the case assumption, $(\hat{\psi}', \hat{K}_{\text{root}}, g'', \hat{\psi}) \in \delta$. By cases on $(\hat{\psi}', [\vec{g'}]) \overset{\hat{K}'', g''}{\underset{G}{\longmapsto}} (\hat{\psi}, \hat{\kappa})$:

Case $(\hat{\psi}', [\vec{g'}]) \overset{\hat{K}'', \hat{\phi}_+}{\underset{G}{\longmapsto}} (\hat{\psi}, \hat{\kappa})$**.**

By definition of $\hat{f}$, $(\hat{\psi}', A) \overset{\hat{\phi}_+}{\rightarrowtail} (\hat{\psi}, A \cup \mathscr{T}(\hat{\phi}_+)) \in G'$. By definition of *StackRoot* and A, $\mathit{StackRoot}([\vec{g'}g'']) = \mathit{StackRoot}([\vec{g}]) = \mathit{StackRoot}(\hat{\kappa})$.

Case $(\hat{\psi}', [\vec{g'}]) \overset{\hat{K}'', \varepsilon}{\underset{G}{\longmapsto}} (\hat{\psi}, \hat{\kappa})$**.**

Similar to previous case.

Case $(\hat{\psi}', [\vec{g'}]) \overset{\hat{K}'', \hat{\phi}_-}{\underset{G}{\longmapsto}} (\hat{\psi}, \hat{\kappa})$**.**

Since is $[\vec{g'}\hat{\phi}_-]$ is defined, there is an i such that $g_i = \hat{\phi}_+$, which is witness to an edge in the trace with that action, $\hat{\psi}_b \overset{\hat{K}''', \hat{\phi}_+}{\underset{G}{\twoheadrightarrow}} \hat{\psi}_e$ By definition of $[_]$, the actions from $\hat{\psi}_e$ to $\hat{\psi}'$ cancel to ε, meaning by IH $(\hat{\psi}_e, A) \rightarrowtail (\hat{\psi}', A) \in H$, and $(\hat{\psi}_b, A') \overset{\hat{\phi}_+}{\rightarrowtail} (\hat{\psi}_e, A) \in E$. Thus the pop edge is added by definition of $\hat{f}'$. The new balanced path $(\hat{\psi}_b, A') \rightarrowtail (\hat{\psi}, A')$ is also added, and extended paths get added with propagation.

Approximation follows by composition with Theorem 12.1.

(G approximates G'): Directly from 18.7. □

The approximate GC has a similar invariant, except the sets of addresses are with respect to the $\hat{t}$ computation.

$$\mathit{inv}_{\hat{\Gamma}} : \mathsf{Exp} \to \widehat{\mathit{System}}_\Gamma \to \mathit{Prop}$$

$$\begin{aligned}\mathit{inv}_{\hat{\Gamma}}(e)(\overbrace{((\hat{P}, \hat{E}), \hat{H})}^{G}) = (\hat{P} = \bigcup \left\{ \{\hat{\psi}, \hat{\psi}'\} : \hat{\psi} \overset{A,g}{\rightarrowtail} \hat{\psi}' \in \hat{E} \right\}) \\ \wedge \forall \hat{\psi}_0 \overset{A,g}{\rightarrowtail} \hat{\psi}_1 \in \hat{E}. \exists (\hat{\psi}_0^\Gamma, A_\Gamma) \overset{g}{\rightarrowtail} (\hat{\psi}_1^\Gamma, A'_\Gamma) \in \hat{E}. (\forall i. \hat{\psi}_i^\Gamma \sqsubseteq \hat{\psi}_i) \wedge A_\Gamma \subseteq A \\ \wedge \forall \hat{\psi}_0 \rightarrowtail \hat{\psi}_1 \in \hat{H}. \exists (\hat{\psi}_0^\Gamma, A_\Gamma) \rightarrowtail (\hat{\psi}_1^\Gamma, A_\Gamma) \in \hat{H}'. \forall i. \hat{\psi}_i^\Gamma \sqsubseteq \hat{\psi}_i \\ \wedge \forall \hat{\psi} \in \hat{P}. \exists (\hat{\psi}^\Gamma, A) \in \hat{P}'. \hat{\psi}^\Gamma \sqsubseteq \hat{\psi} \wedge \mathrm{lfp}(\hat{t})(\hat{\psi}) \subseteq A\end{aligned}$$

where $((\hat{P}', \hat{E}'), \hat{H}') = \mathrm{lfp}(\hat{f}_e)$

Lemma 18.8 (*Approx GC invariant*)
For all e, if $\mathit{inv}_{\hat{\Gamma}}(e)(G)$ then $\mathit{inv}_{\hat{\Gamma}}(e)(\hat{f}'_e(G))$

Proof
Straightforward case analysis. □

Proof of Theorem 12.3

Proof
Induct on path in $\mathrm{lfp}(\hat{f}_e)$ and apply 18.8. □

19 Haskell implementation of CRPDSs

Where it is critical to understanding the details of the analysis, we have transliterated the formalism into Haskell. We make use of a two extensions in GHC:

```
-XTypeOperators -XTypeSynonymInstances
```

All code is in the context of the following header, and we'll assume the standard instances of type classes like `Ord` and `Eq`.

```
import Prelude hiding ((!!))

import Data.Map as Map hiding (map,foldr)
import Data.Set as Set hiding (map,foldr)
import Data.List as List hiding ((!!))

type ℙ s = Set.Set s
type k :-> v = Map k v

(==>) :: a -> b -> (a,b)
(==>) x y = (x,y)

(//) :: Ord a => (a :-> b) -> [(a,b)] -> (a :-> b)
(//) f [(x,y)] = Map.insert x y f

set x = Set.singleton x
```

19.1 Transliteration of NFA formalism

We represent an NFA as a set of labeled forward edges, the inverse of those edges (for convenience), a start state and an end state:

```
type NFA state char =
 (NFAEdges state char,NFAEdges state char,state,state)
type NFAEdges state char = state :-> ℙ(Maybe char,state)
```

19.2 ANF

```
data Exp    = Ret AExp
            | App Call
            | Let1 Var Call Exp
data AExp   = Ref Var
            | Lam Lambda
data Lambda = Var :=> Exp
data Call   = AExp :@ AExp
type Var    = String
```

```
-- Abstract state-space:
type AConf  = (Exp,AEnv,AStore,AKont)
type AEnv   = Var :-> AAddr
type AStore = AAddr :-> AD
type AD     =  (AVal)
data AVal   = AClo (Lambda, AEnv)
type AKont  = [AFrame]
type AFrame = (Var,Exp,AEnv)

data AAddr = ABind Var AContext
type AContext = [Call]
```

Abstract configuration space transliterated into Haskell. In the code, we defined abstract addresses to be able to support k-CFA-style polyvariance.

Atomic expression evaluation implementation:

```
aeval :: (AExp,AEnv,AStore) -> AD
aeval (Ref v, ρ, σ) = σ!!(ρ!v)
aeval (Lam l, ρ, σ) = set $ AClo (l, ρ)
```

We encode the transition relation it as a function that returns lists of states:

```
 astep :: AConf -> [AConf]

astep (App (f :@ ae), ρ, σ, κ) = [(e, ρ'', σ', κ) |
   AClo(v :=> e, ρ') <- Set.toList $ aeval(f, ρ, σ),
   let a = aalloc(v, App (f :@ ae)),
   let ρ'' = ρ' // [v ==> a],
   let σ' = σ  [a ==> aeval(ae, ρ, σ)] ]
astep (Let1 v call e, ρ, σ, κ) =
  [(App call, ρ, σ, (v, e, ρ) : κ)]
astep (Ret ae, ρ, σ, (v, e, ρ') : κ) = [(e, ρ'', σ', κ)]
 where a = aalloc(v, Ret ae)
       ρ'' = ρ' // [v ==> a]
       σ' = σ  [a ==> aeval(ae, ρ, σ)]
```

19.3 Partial orders

We define a typeclass for lattices:

```
class Lattice a where
 bot :: a
 top :: a
 (⊑) :: a -> a -> Bool
 (⊔) :: a -> a -> a
 (⊓) :: a -> a -> a
```

And, we can lift instances to sets and maps:

```
instance (Ord s, Eq s) => Lattice (ℙ s) where
 bot = Set.empty
 top = error "no representation of universal set"
 x ⊔ y = x `Set.union` y
 x ⊓ y = x `Set.intersection` y
 x ⊑ y = x `Set.isSubsetOf` y

instance (Ord k, Lattice v) => Lattice (k :-> v) where
 bot = Map.empty
 top = error "no representation of top map"
 f ⊑ g = Map.isSubmapOfBy (⊑) f g
 f ⊔ g = Map.unionWih (⊔) f g
 f ⊓ g = Map.intersectionWith (⊓) f g

(⊔) :: (Ord k, Lattice v) => (k :-> v) -> [(k,v)] -> (k :-> v)
f ⊔ [(k,v)] = Map.insertWith (⊔) k v f

(!!) :: (Ord k, Lattice v) => (k :-> v) -> k -> v
f !! k = Map.findWithDefault bot k f
```

19.4 Reachability

We can turn any data type to a stack-action alphabet:

```
data StackAct frame = Push { frame :: frame }
                    | Pop  { frame :: frame }
                    | Unch

type CRPDS control frame = (Edges control frame, control)
type Edges control frame = control :->  (StackAct frame,control)
```

We split the encoding of δ into two functions for efficiency purposes:

```
type Delta control frame =
 (TopDelta control frame, NopDelta control frame)
type TopDelta control frame =
 control -> frame -> [(control,StackAct frame)]
type NopDelta control frame =
 control -> [(control,StackAct frame)]
```

If we only want to know push and no-change transitions, we can find these with a `NopDelta` function without providing the frame that is currently on top of the stack. If we want pop transitions as well, we can find these with a `TopDelta` function, but of course, it must have

access to the top of the stack. In practice, a `TopDelta` function would suffice, but there are situations where only push and no-change transitions are needed, and having access to `NopDelta` avoids extra computation.

At this point, we must clarify how to embed the abstract transition relation into a pushdown transition relation:

```
adelta :: TopDelta AControl AFrame
adelta (e, ρ, σ) γ = [ ((e', ρ', σ'), g) |
 (e', ρ', σ', κ) <- astep (e, ρ, σ, [γ]),
 let g = case κ of
          []        -> Pop γ
          [γ1 , _ ] -> Push γ1
          [ _ ]     -> Unch ]

adelta' :: NopDelta AControl AFrame
adelta' (e, ρ, σ) = [ ((e', ρ', σ'), g) |
 (e', ρ', σ', κ) <- astep (e, ρ, σ, []),
 let g = case κ of
          [γ1]  -> Push γ1
          [ ]   -> Unch ]
```

The function `crpds` will invoke the fixed point solver:

```
crpds :: (Ord control, Ord frame) =>
       (Delta control frame) ->
       control ->
       frame ->
       CRPDS control frame
crpds (δ,δ') q0 0 =
 (summarize (δ,δ') etg1 ecg1 [] dE dH, q0) where
 etg1 = (Map.empty // [q0 ==> Set.empty],
         Map.empty // [q0 ==> Set.empty])
 ecg1 = (Map.empty // [q0 ==> set q0],
         Map.empty // [q0 ==> set q0])
 (dE,dH) = sprout (δ,δ') q0
```

Figure 11 provides the code for `summarize`, which conducts the fixed point calculation, the executable equivalent of Figure 6:

```
summarize :: (Ord control, Ord frame) =>
             (Delta control frame) ->
             (ETG control frame) ->
             (ECG control) ->
             [control] ->
             [Edge control frame] ->
             [EpsEdge control] ->
             (Edges control frame)
```

To expose the structure of the computation, we've added a few types:

```
-- A set of edges, encoded as a map:
type Edges control frame =
 control :-> ℙ (StackAct frame,control)

-- Epsilon edges:
type EpsEdge control = (control,control)

-- Explicit transition graph:
type ETG control frame =
 (Edges control frame, Edges control frame)

-- Epsilon closure graph:
type ECG control =
 (control :-> ℙ(control), control :-> ℙ(control))
```

An explicit transition graph is an explicit encoding of the reachable subset of the transition relation. The function `summarize` takes six parameters:

1. the pushdown transition function;
2. the current explicit transition graph;
3. the current ε-closure graph;
4. a work-list of states to add;
5. a work-list of explicit transition edges to add; and
6. a work-list of ε-closure transition edges to add.

The function `summarize` processes ε-closure edges first, then explicit transition edges and then individual states. It *must* process ε-closure edges first to ensure that the ε-closure graph is closed when considering the implications of other edges.

Sprouting

```
summarize (δ,δ') (fw,bw) (fe,be) [] [] [] = fw

summarize (δ,δ') (fw,bw) (fe,be) (q:dS) [] []
 | fe `contains` q = summarize (δ,δ') (fw,bw) (fe,be) dS [] []
summarize (δ,δ') (fw,bw) (fe,be) (q:dS) [] [] =
 summarize (δ,δ') (fw',bw') (fe',be') dS dE' dH' where
  (dE',dH') = sprout (δ,δ') q
  fw' = fw ⊔ [q ==> Set.empty]
  bw' = bw ⊔ [q ==> Set.empty]
  fe' = fe ⊔ [q ==> set q]
  be' = be ⊔ [q ==> set q]

summarize (δ,δ') (fw,bw) (fe,be) dS ((q,g,q'):dE) []
 | (q,g,q') `isin` fw  = summarize (δ,δ') (fw,bw) (fe,be) dS dE []
summarize (δ,δ') (fw,bw) (fe,be) dS ((q,Push ,q'):dE) [] =
 summarize (δ,δ') (fw',bw') (fe',be') dS' dE'' dH' where
  (dE',dH') = addPush (fw,bw) (fe,be) (δ,δ') (q,Push ,q')
  dE'' = dE' ++ dE''
  dS' = q':dS
  fw' = fw ⊔ [q  ==> set (Push ,q')]
  bw' = bw ⊔ [q' ==> set (Push ,q) ]
  fe' = fe ⊔ [q  ==> set q ]
  be' = fe ⊔ [q' ==> set q']

summarize (δ,δ') (fw,bw) (fe,be) dS ((q,Pop ,q'):dE) [] =
 summarize (δ,δ') (fw',bw') (fe',be') dS' dE'' dH' where
  (dE',dH') = addPop (fw,bw) (fe,be) (δ,δ') (q,Pop ,q')
  dE'' = dE ++ dE'
  dS' = q':dS
  fw' = fw ⊔ [q  ==> set (Pop ,q')]
  bw' = bw ⊔ [q' ==> set (Pop ,q) ]
  fe' = fe ⊔ [q  ==> set q ]
  be' = fe ⊔ [q' ==> set q']

summarize (δ,δ') (fw,bw) (fe,be) dS ((q,Unch,q'):dE) [] =
 summarize (δ,δ') (fw',bw') (fe',be') dS' dE [(q,q')] where
  dS' = q':dS
  fw' = fw ⊔ [q  ==> set (Unch,q')]
  bw' = bw ⊔ [q' ==> set (Unch,q) ]
  fe' = fe ⊔ [q  ==> set q ]
  be' = fe ⊔ [q' ==> set q']

summarize (δ,δ') (fw,bw) (fe,be) dS dE ((q,q'):dH)
 | (q,q') `isin` fe  = summarize (δ,δ') (fw,bw) (fe,be) dS dE dH
summarize (δ,δ') (fw,bw) (fe,be) dS dE ((q,q'):dH) =
 summarize (δ,δ') (fw,bw) (fe',be') dS dE' dH' where
   (dE',dH') = addEmpty (fw,bw) (fe,be) (δ,δ') (q,q')
   fe' = fe ⊔ [q  ==> set q ]
   be' = fe ⊔ [q' ==> set q']
```

Fig. 11: An implementation of pushdown control-state reachability.

```
sprout :: (Ord control) =>
          Delta control frame ->
          control ->
          ([Edge control frame], [EpsEdge control])
sprout (δ,δ') q = (dE, dH) where
  edges = δ' q
  dE = [ (q,g,q') | (q',g) <- edges, isPush g ]
  dH = [ (q,q')   | (q',g) <- edges, isUnch g ]
```

Pushing

```
addPush :: (Ord control) =>
           ETG control frame ->
           ECG control ->
           Delta control frame ->
           Edge control frame ->
           ([Edge control frame], [EpsEdge control])
addPush (fw,bw) (fe,be) (δ,δ') (s,Push γ,q) = (dE,dH) where
 qset' = Set.toList $ fe!q
 dE = [ (q',g,q'') | q' <- qset', (q'',g) <- δ q' γ, isPop g ]
 dH = [ (s,q'')    | (q',Pop _,q'') <- dE ]
```

Popping

```
addPop :: (Ord control) =>
          ETG control frame ->
          ECG control ->
          Delta control frame ->
          Edge control frame ->
          ([Edge control frame], [EpsEdge control])
addPop (fw,bw) (fe,be) (δ,δ') (s'',Pop γ,q) = (dE,dH) where
 sset' = Set.toList $ be!s''
 dH = [ (s,q) | s' <- sset',
                (g,s) <- Set.toList $ bw!s', isPush g ]
 dE = []
```

Clearly, we could eliminate the new edges parameter `dE` for the function `addPop`, but we have retained it for stylistic symmetry.

Adding empty edges The function `addEmpty` has many cases to consider:

```
addEmpty :: (Ord control) =>
            ETG control frame ->
            ECG control ->
            Delta control frame ->
            EpsEdge control ->
            ([Edge control frame], [EpsEdge control])
addEmpty (fw,bw) (fe,be) (δ,δ') (s'',s''') = (dE,dH) where
 sset'    = Set.toList $ be!s''
 sset'''' = Set.toList $ fe!s'''
 dH'   = [ (s',s'''')  | s' <- sset', s'''' <- sset'''' ]
 dH''  = [ (s',s''')   | s' <- sset' ]
 dH''' = [ (s'',s'''') | s'''' <- sset'''' ]

 sEdges = [ (g,s) | s' <- sset', (g,s) <- Set.toList $ bw!s' ]

 dE = [ (s'''',g',q) | s'''' <- sset'''',
                       (g,s) <- sEdges,
                       isPush g, let Push γ = g,
                       (q,g') <- δ s'''' γ,
                       isPop g' ]

 dH'''' = [ (s,q) | (_,s) <- sEdges, (_,_,q) <- dE ]

 dH = dH' ++ dH'' ++ dH''' ++ dH''''
```